



US010002046B2

(12) **United States Patent**
Cohen

(10) **Patent No.:** **US 10,002,046 B2**
(45) **Date of Patent:** ***Jun. 19, 2018**

(54) **ZERO-ONE BALANCE MANAGEMENT IN A SOLID-STATE DISK CONTROLLER**

(71) Applicant: **Seagate Technology LLC**, Cupertino, CA (US)

(72) Inventor: **Earl T. Cohen**, Oakland, CA (US)

(73) Assignee: **Seagate Technology LLC**, Cupertino, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days. days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **15/238,839**

(22) Filed: **Aug. 17, 2016**

(65) **Prior Publication Data**

US 2016/0357631 A1 Dec. 8, 2016

Related U.S. Application Data

(63) Continuation of application No. 14/472,869, filed on Aug. 29, 2014, now Pat. No. 9,430,154, which is a (Continued)

(51) **Int. Cl.**
G06F 11/10 (2006.01)
H03M 13/11 (2006.01)
(Continued)

(52) **U.S. Cl.**
CPC **G06F 11/1068** (2013.01); **G06F 3/061** (2013.01); **G06F 3/0619** (2013.01);
(Continued)

(58) **Field of Classification Search**
CPC **G06F 3/061**; **G06F 3/0619**; **G06F 3/0655**;
G06F 3/0676; **G06F 3/0679**;
(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,341,085 B1 1/2002 Yamagami et al.
6,347,051 B2 2/2002 Yamagami et al.

(Continued)

FOREIGN PATENT DOCUMENTS

JP 2009163774 7/2009
KR 1020100004731 1/2010

(Continued)

OTHER PUBLICATIONS

Cohen, Earl T.; Final Office Action for U.S. Appl. No. 14/472,869, filed Aug. 29, 2014, dated Mar. 7, 2016, 28 pgs.

(Continued)

Primary Examiner — April Y Blair

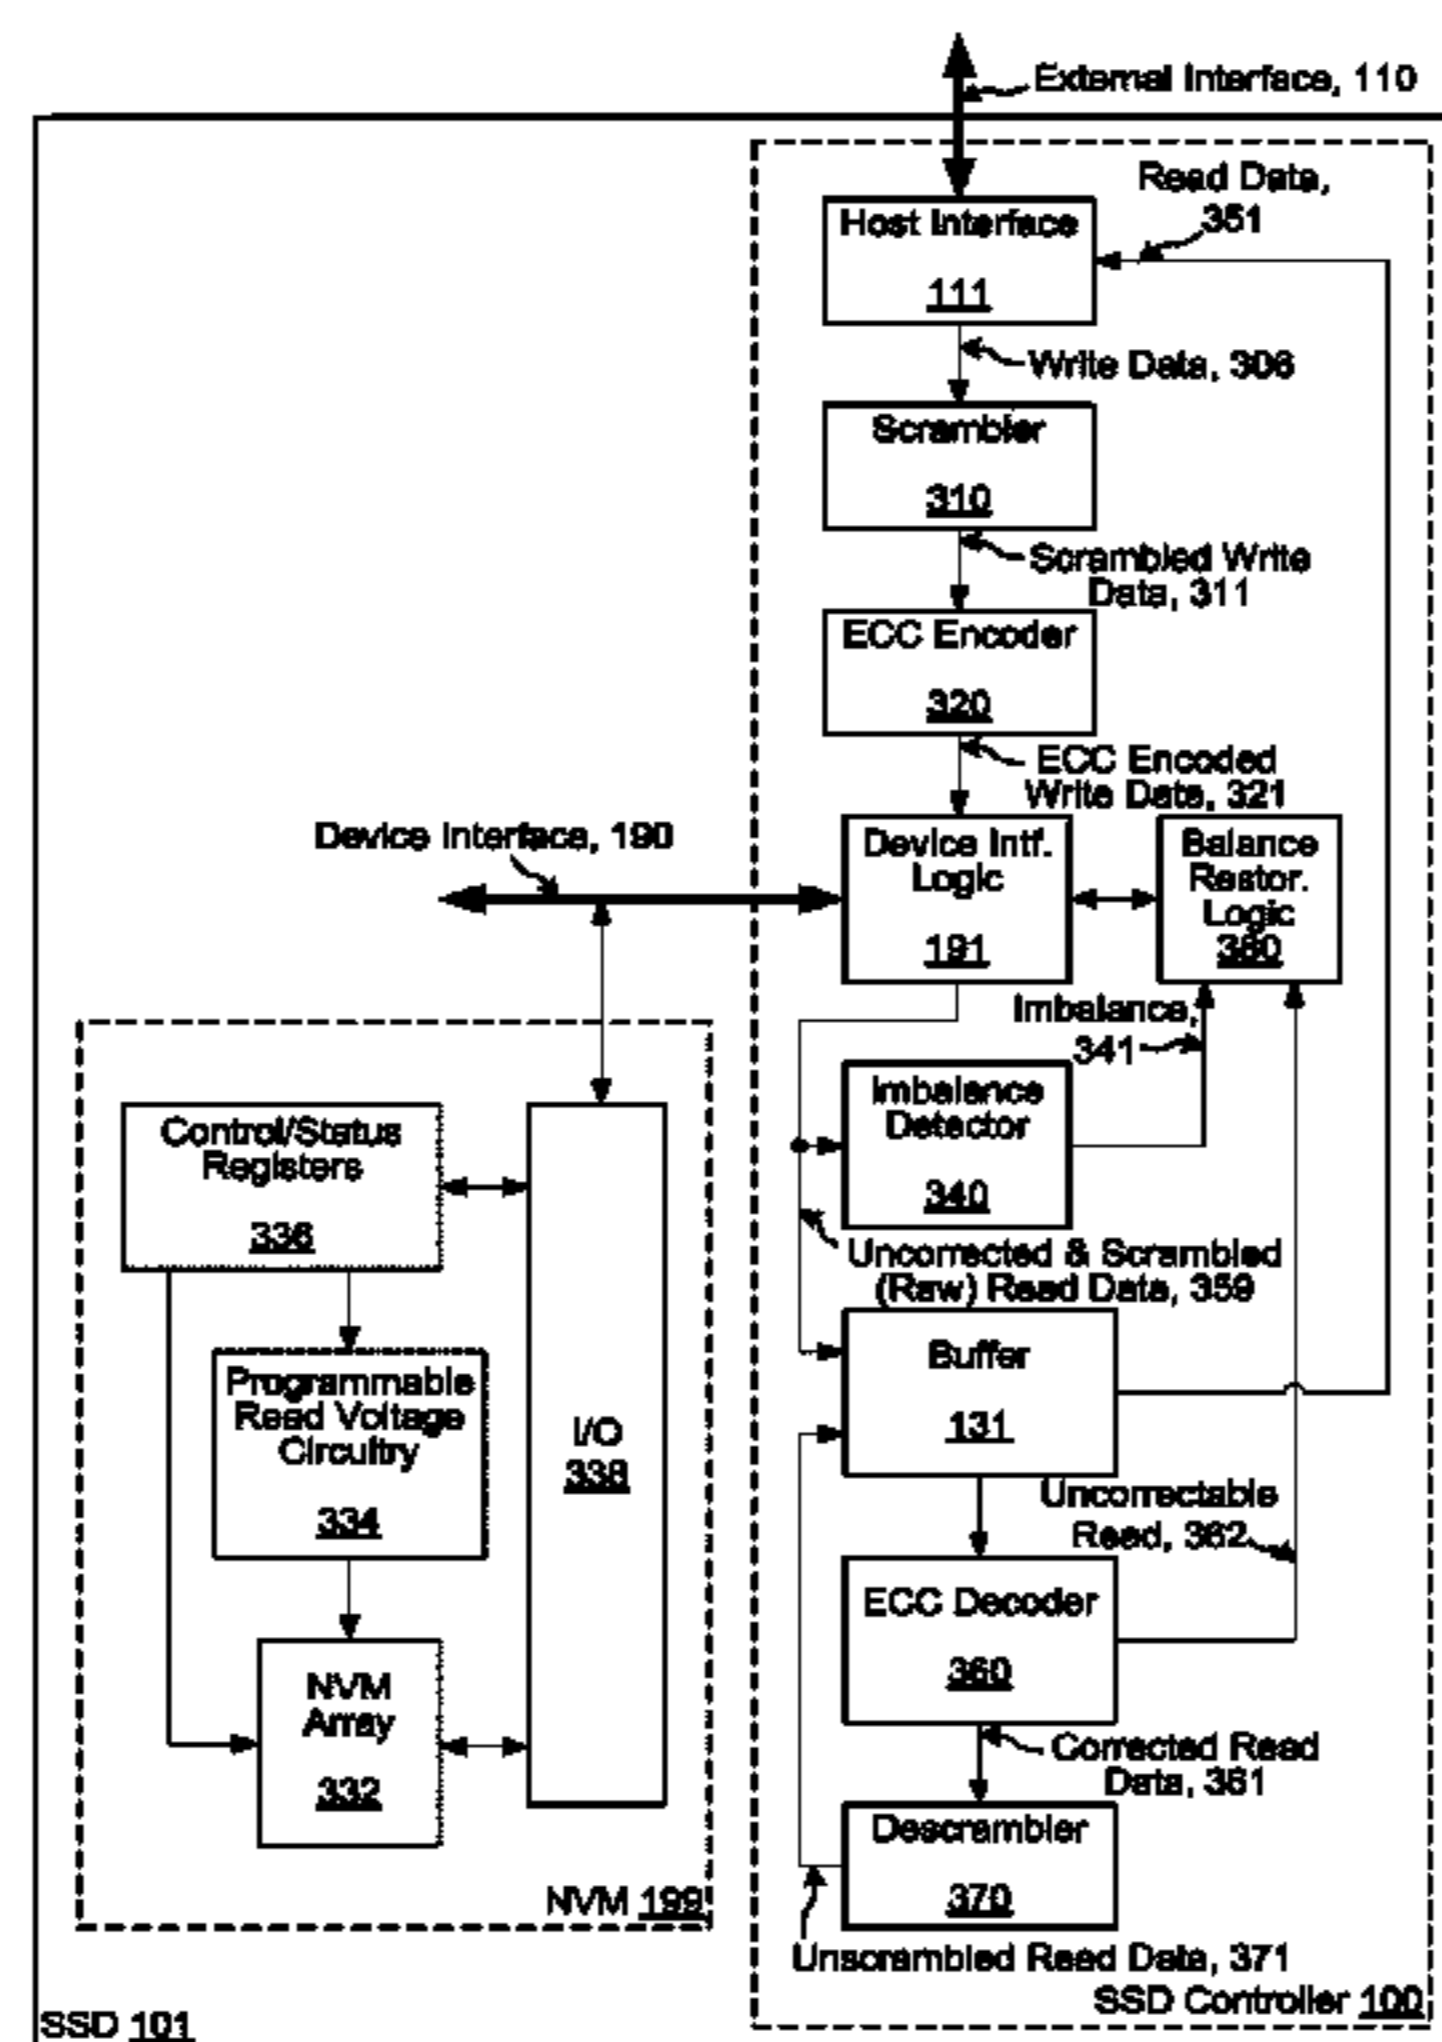
Assistant Examiner — Sazzad Hossain

(74) *Attorney, Agent, or Firm* — Taylor English Duma LLP

(57) **ABSTRACT**

An SSD controller maintains a zero count and a one count, and/or in some embodiments a zero/one disparity count, for each read unit read from an SLC NVM (or the lower pages of an MLC). In an event that the read unit is uncorrectable in part due to a shift in the threshold voltage distributions away from their nominal distributions, the maintained counts enable a determination of a direction and/or a magnitude to adjust a read threshold to track the threshold voltage shift and restore the read data zero/one balance. In various embodiments, the adjusted read threshold is determined in a variety of described ways (counts, percentages) that are based on a number of described factors (determined threshold voltage distributions, known stored values, past NVM operating events). Extensions of the forgoing techniques are described for MLC memories.

20 Claims, 6 Drawing Sheets



Related U.S. Application Data

continuation of application No. 13/464,433, filed on May 4, 2012, now Pat. No. 8,839,073.

(51) **Int. Cl.**

G06F 3/06 (2006.01)
G11C 29/52 (2006.01)
H03M 13/09 (2006.01)
H03M 13/15 (2006.01)
H03M 13/19 (2006.01)

(52) **U.S. Cl.**

CPC **G06F 3/0655** (2013.01); **G06F 3/0658** (2013.01); **G06F 3/0676** (2013.01); **G06F 3/0679** (2013.01); **G06F 11/1012** (2013.01); **G06F 11/1048** (2013.01); **G06F 11/1072** (2013.01); **G11C 29/52** (2013.01); **H03M 13/1102** (2013.01); **G06F 2206/1014** (2013.01); **H03M 13/09** (2013.01); **H03M 13/152** (2013.01); **H03M 13/1515** (2013.01); **H03M 13/19** (2013.01)

(58) **Field of Classification Search**

CPC G06F 11/1012; G06F 11/1048; G06F 2206/1014; H03M 13/09; H03M 13/1102; H03M 13/1515; H03M 13/152; H03M 13/19

See application file for complete search history.

(56)

References Cited

U.S. PATENT DOCUMENTS

6,678,823	B1	1/2004	Fernandez et al.
7,457,178	B2	11/2008	Tu et al.
7,460,402	B2	12/2008	Shibata
7,791,938	B2	9/2010	Kang et al.
7,936,601	B2	5/2011	Kang et al.
8,000,135	B1	8/2011	Perlmutter
8,040,725	B2	10/2011	Kang
8,089,813	B2	1/2012	Dreps
8,125,827	B2	2/2012	Park
8,130,544	B2	3/2012	Chou et al.
8,261,157	B2	9/2012	Litsyn et al.
8,305,812	B2	11/2012	Levy et al.
8,345,477	B1	1/2013	Yang
8,351,258	B1	1/2013	Yang et al.
8,402,349	B2	3/2013	Lee et al.
8,453,022	B2	5/2013	Katz
8,456,926	B2	6/2013	Ong et al.
8,531,888	B2	9/2013	Chilappagari et al.
8,839,073	B2	9/2014	Cohen
9,430,154	B2	8/2016	Cohen
2001/0030890	A1	10/2001	Yamagami et al.
2006/0092772	A1*	5/2006	Tsai G11B 20/10 369/30.01
2007/0159891	A1	7/2007	Tu et al.
2008/0151618	A1	6/2008	Sharon et al.
2008/0175063	A1	7/2008	Son et al.
2008/0263265	A1	10/2008	Litsyn
2009/0003057	A1	1/2009	Kang et al.
2009/0168543	A1	7/2009	Kim et al.
2009/0282186	A1	11/2009	Mokhlesi
2010/0002506	A1	1/2010	Cho et al.
2010/0008151	A1	1/2010	Hwang
2010/0020611	A1	1/2010	Park
2010/0034021	A1	2/2010	Joo
2010/0091535	A1	4/2010	Sommer et al.

2010/0131809	A1	5/2010	Katz
2010/0195835	A1*	8/2010	Nociolo H04L 7/00 380/287
2010/0214853	A1	8/2010	Kang
2010/0277979	A1	11/2010	Kang et al.
2010/0277983	A1	11/2010	Mokhlesi et al.
2011/0032749	A1	2/2011	Liu et al.
2011/0032760	A1	2/2011	Takagiwa
2011/0157992	A1	6/2011	Strasser et al.
2011/0182119	A1*	7/2011	Strasser G11C 16/28 365/185.03
2011/0216603	A1	9/2011	Han et al.
2012/0005558	A1	1/2012	Steiner et al.
2012/0183003	A1*	7/2012	Hwang H03M 13/2732 370/474
2013/0297988	A1	11/2013	Wu et al.
2013/0343131	A1	12/2013	Wu et al.
2014/0095110	A1	4/2014	Chen et al.
2015/0205527	A1	7/2015	Cohen
2015/0287453	A1	10/2015	Wu

FOREIGN PATENT DOCUMENTS

KR	1020100006344	1/2010
KR	1020100010746	2/2010
KR	10-1005188	1/2011
TW	200929238	7/2009
TW	200949850	12/2009
TW	201035975	10/2010
WO	2011113034	9/2011
WO	2013166200	11/2013

OTHER PUBLICATIONS

Cohen, Earl T.; Issue Notification for U.S. Appl. No. 14/472,869, filed Aug. 29, 2014, dated Aug. 10, 2016, 1 pg.

Cohen, Earl T.; Non-Final Office Action for U.S. Appl. No. 14/472,869, filed Aug. 29, 2014, dated Jul. 2, 2015, 20 pgs.

Cohen, Earl T.; Notice of Allowance for U.S. Appl. No. 14/472,869, filed Aug. 29, 2014, dated Apr. 15, 2016, 9 pgs.

Cohen, Earl T.; Notice of Allowance for U.S. Appl. No. 14/472,869, filed Aug. 29, 2014, dated May 4, 2016, 17 pgs.

Cohen, Earl T.; European Office Action for serial No. 13784892.5, filed May 1, 2013, dated Sep. 29, 2016, 6 pgs.

Cohen, Earl T.; International Preliminary Report on Patentability for PCT/US2013/039138, filed May 1, 2013, dated Nov. 4, 2014, 6 pgs.

Cohen, Earl T.; International Search Report and Written Opinion for PCT/US2013/039138, filed May 1, 2013, dated Aug. 23, 2013, 8 pgs.

LSI Corporation; International Search Report and Written Opinion for PCT Application No. PCT/US2013/046422, Filed Jun. 18, 2013, dated Oct. 22, 2013, 8 pgs.

LSI Corporation; Taiwan Office Action for serial No. 102115758, filed May 2, 2013, dated May 6, 2015, 10 pgs.

Seagate Technology, LLC; Extended European Search Report for serial No. 13784892.5, filed May 1, 2013, dated Oct. 23, 2015, 6 pgs.

Seagate Technology, LLC; Japanese Office Action for serial No. 2015-510441, filed May 1, 2013, dated Jan. 24, 2017, 4 pgs.

LSI Corporation; Final Office Action for Taiwanese application No. 102115758, filed May 2, 2013, dated Mar. 2, 2017, 5 pgs.

Seagate Technology, LLC; Notice of Allowance for Korean application No. 10-2014-7032045, filed May 1, 2013, dated Mar. 30, 2017, 3 pgs.

Seagate Technology LLC; Notice of Grant for Japanese serial No. 2015-510441, filed May 1, 2013, dated Jun. 13, 2017, 3 pgs.

* cited by examiner

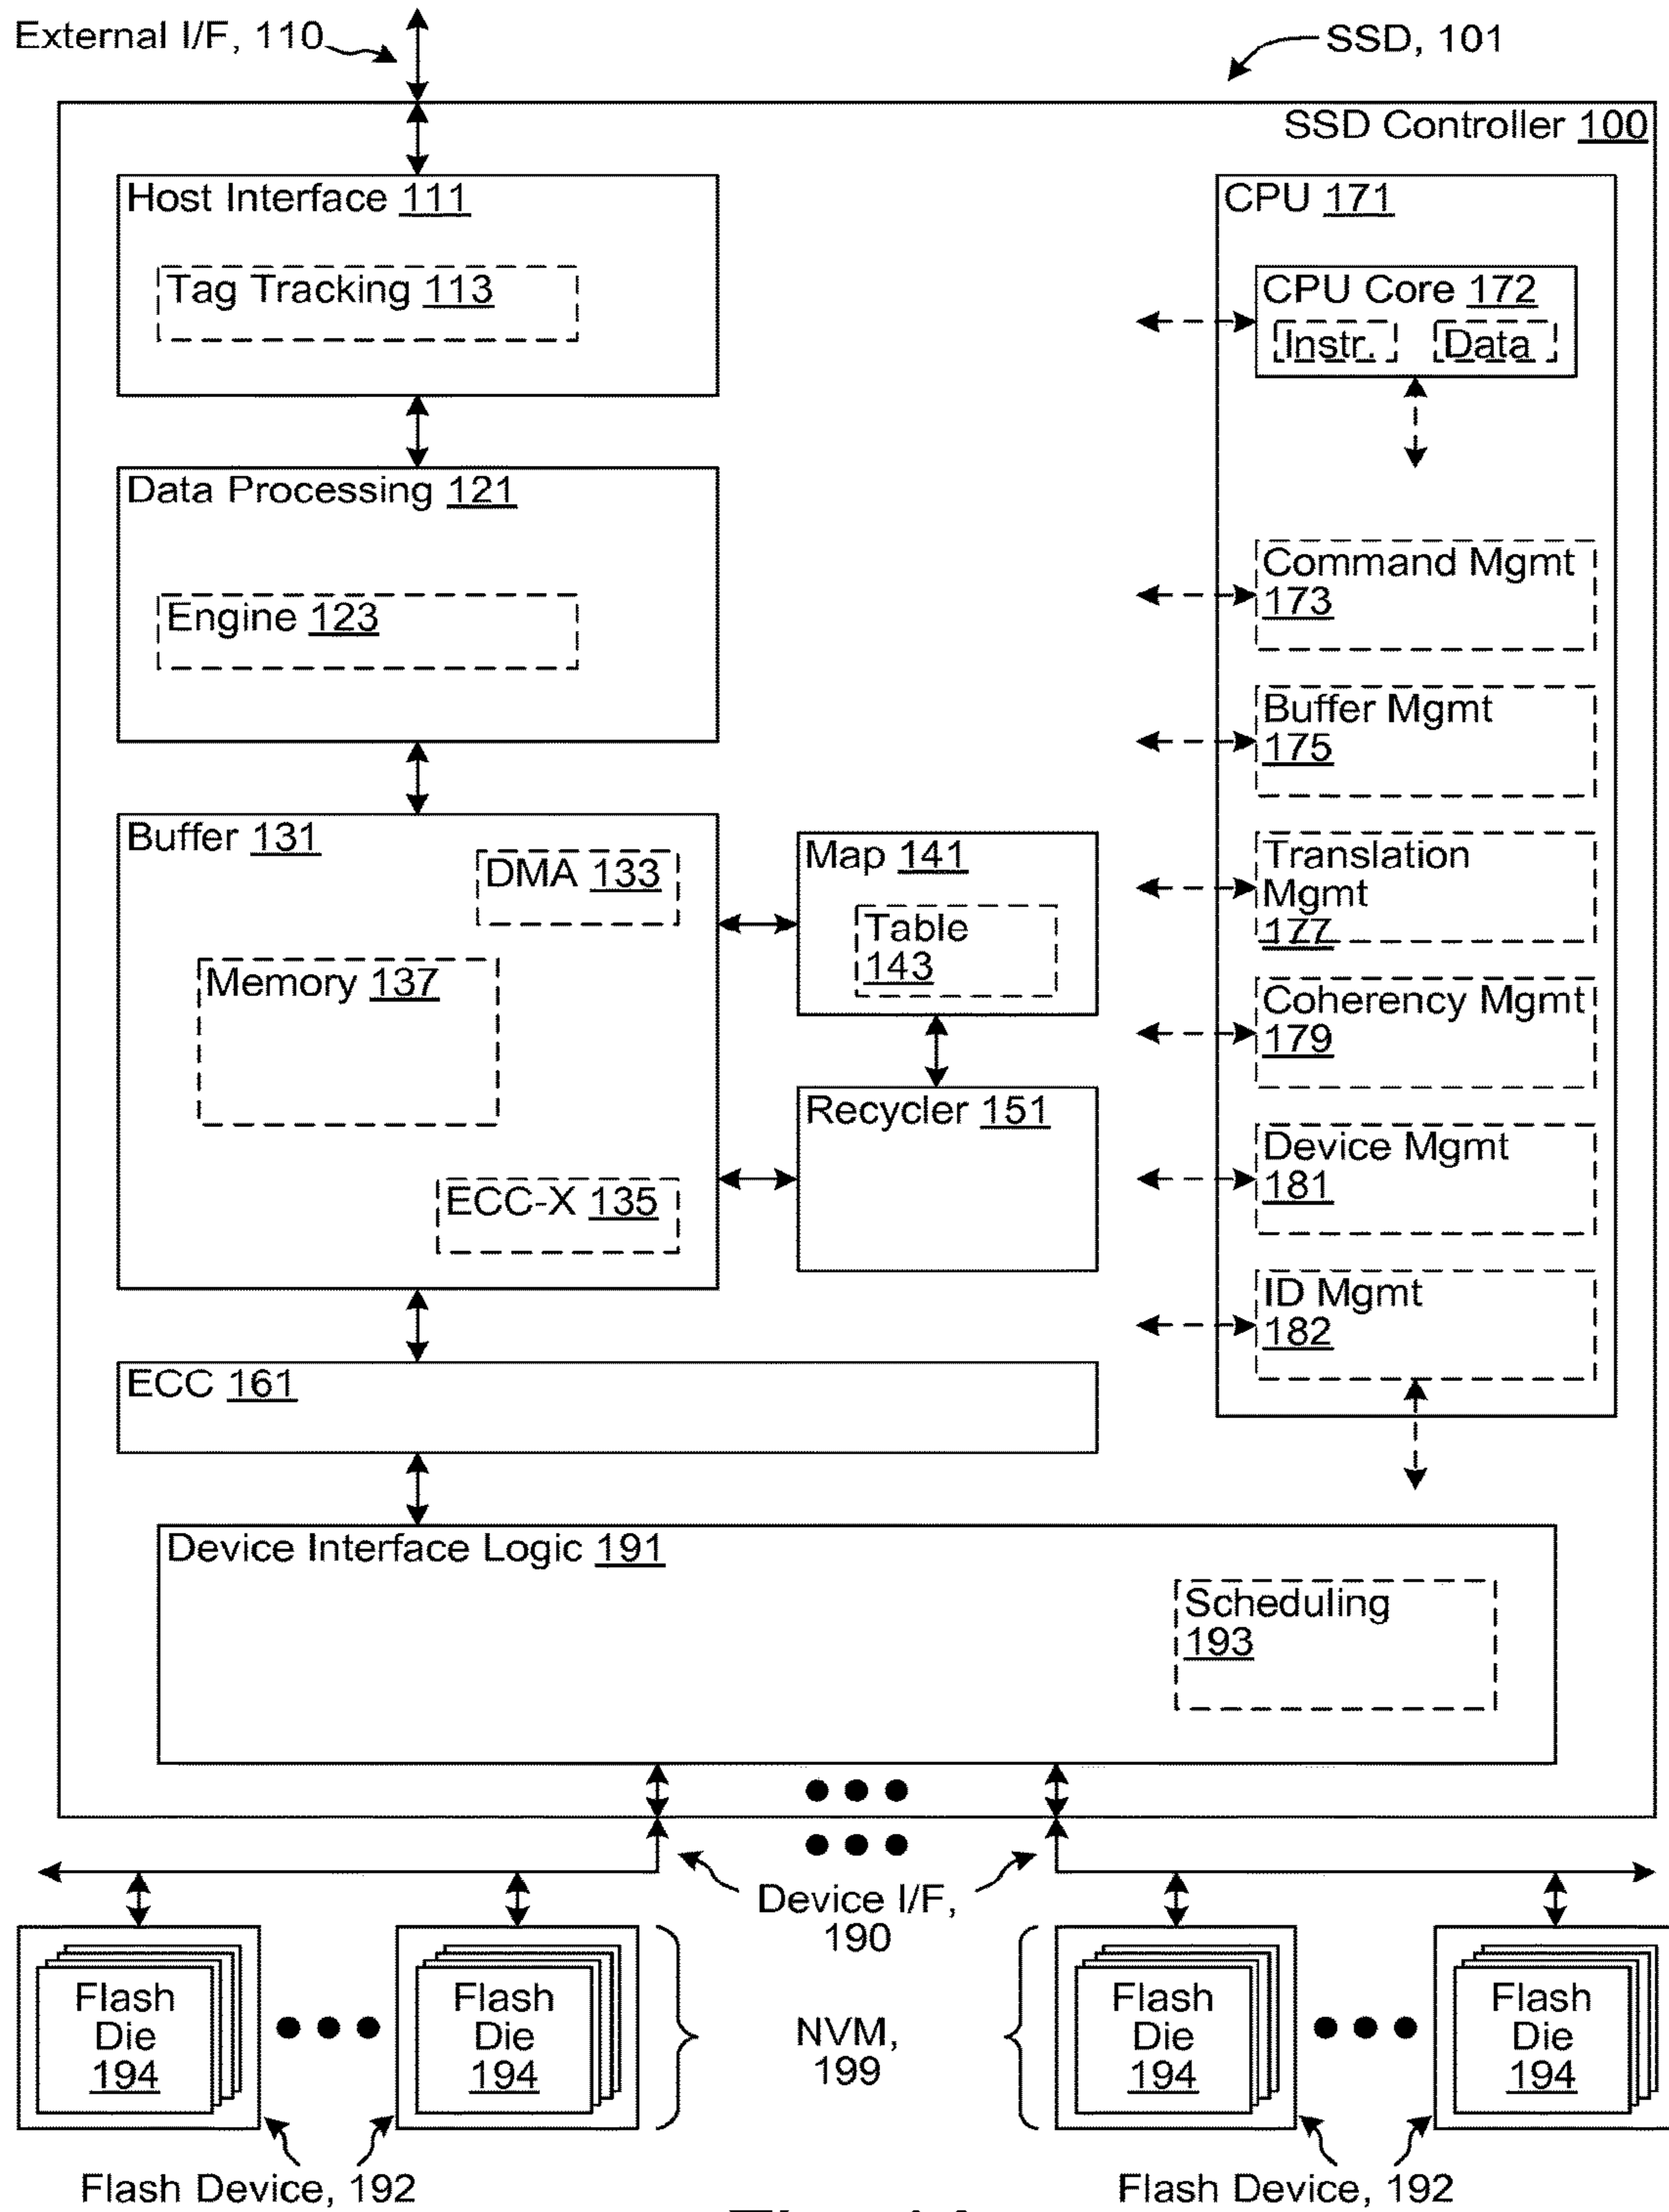


Fig. 1A

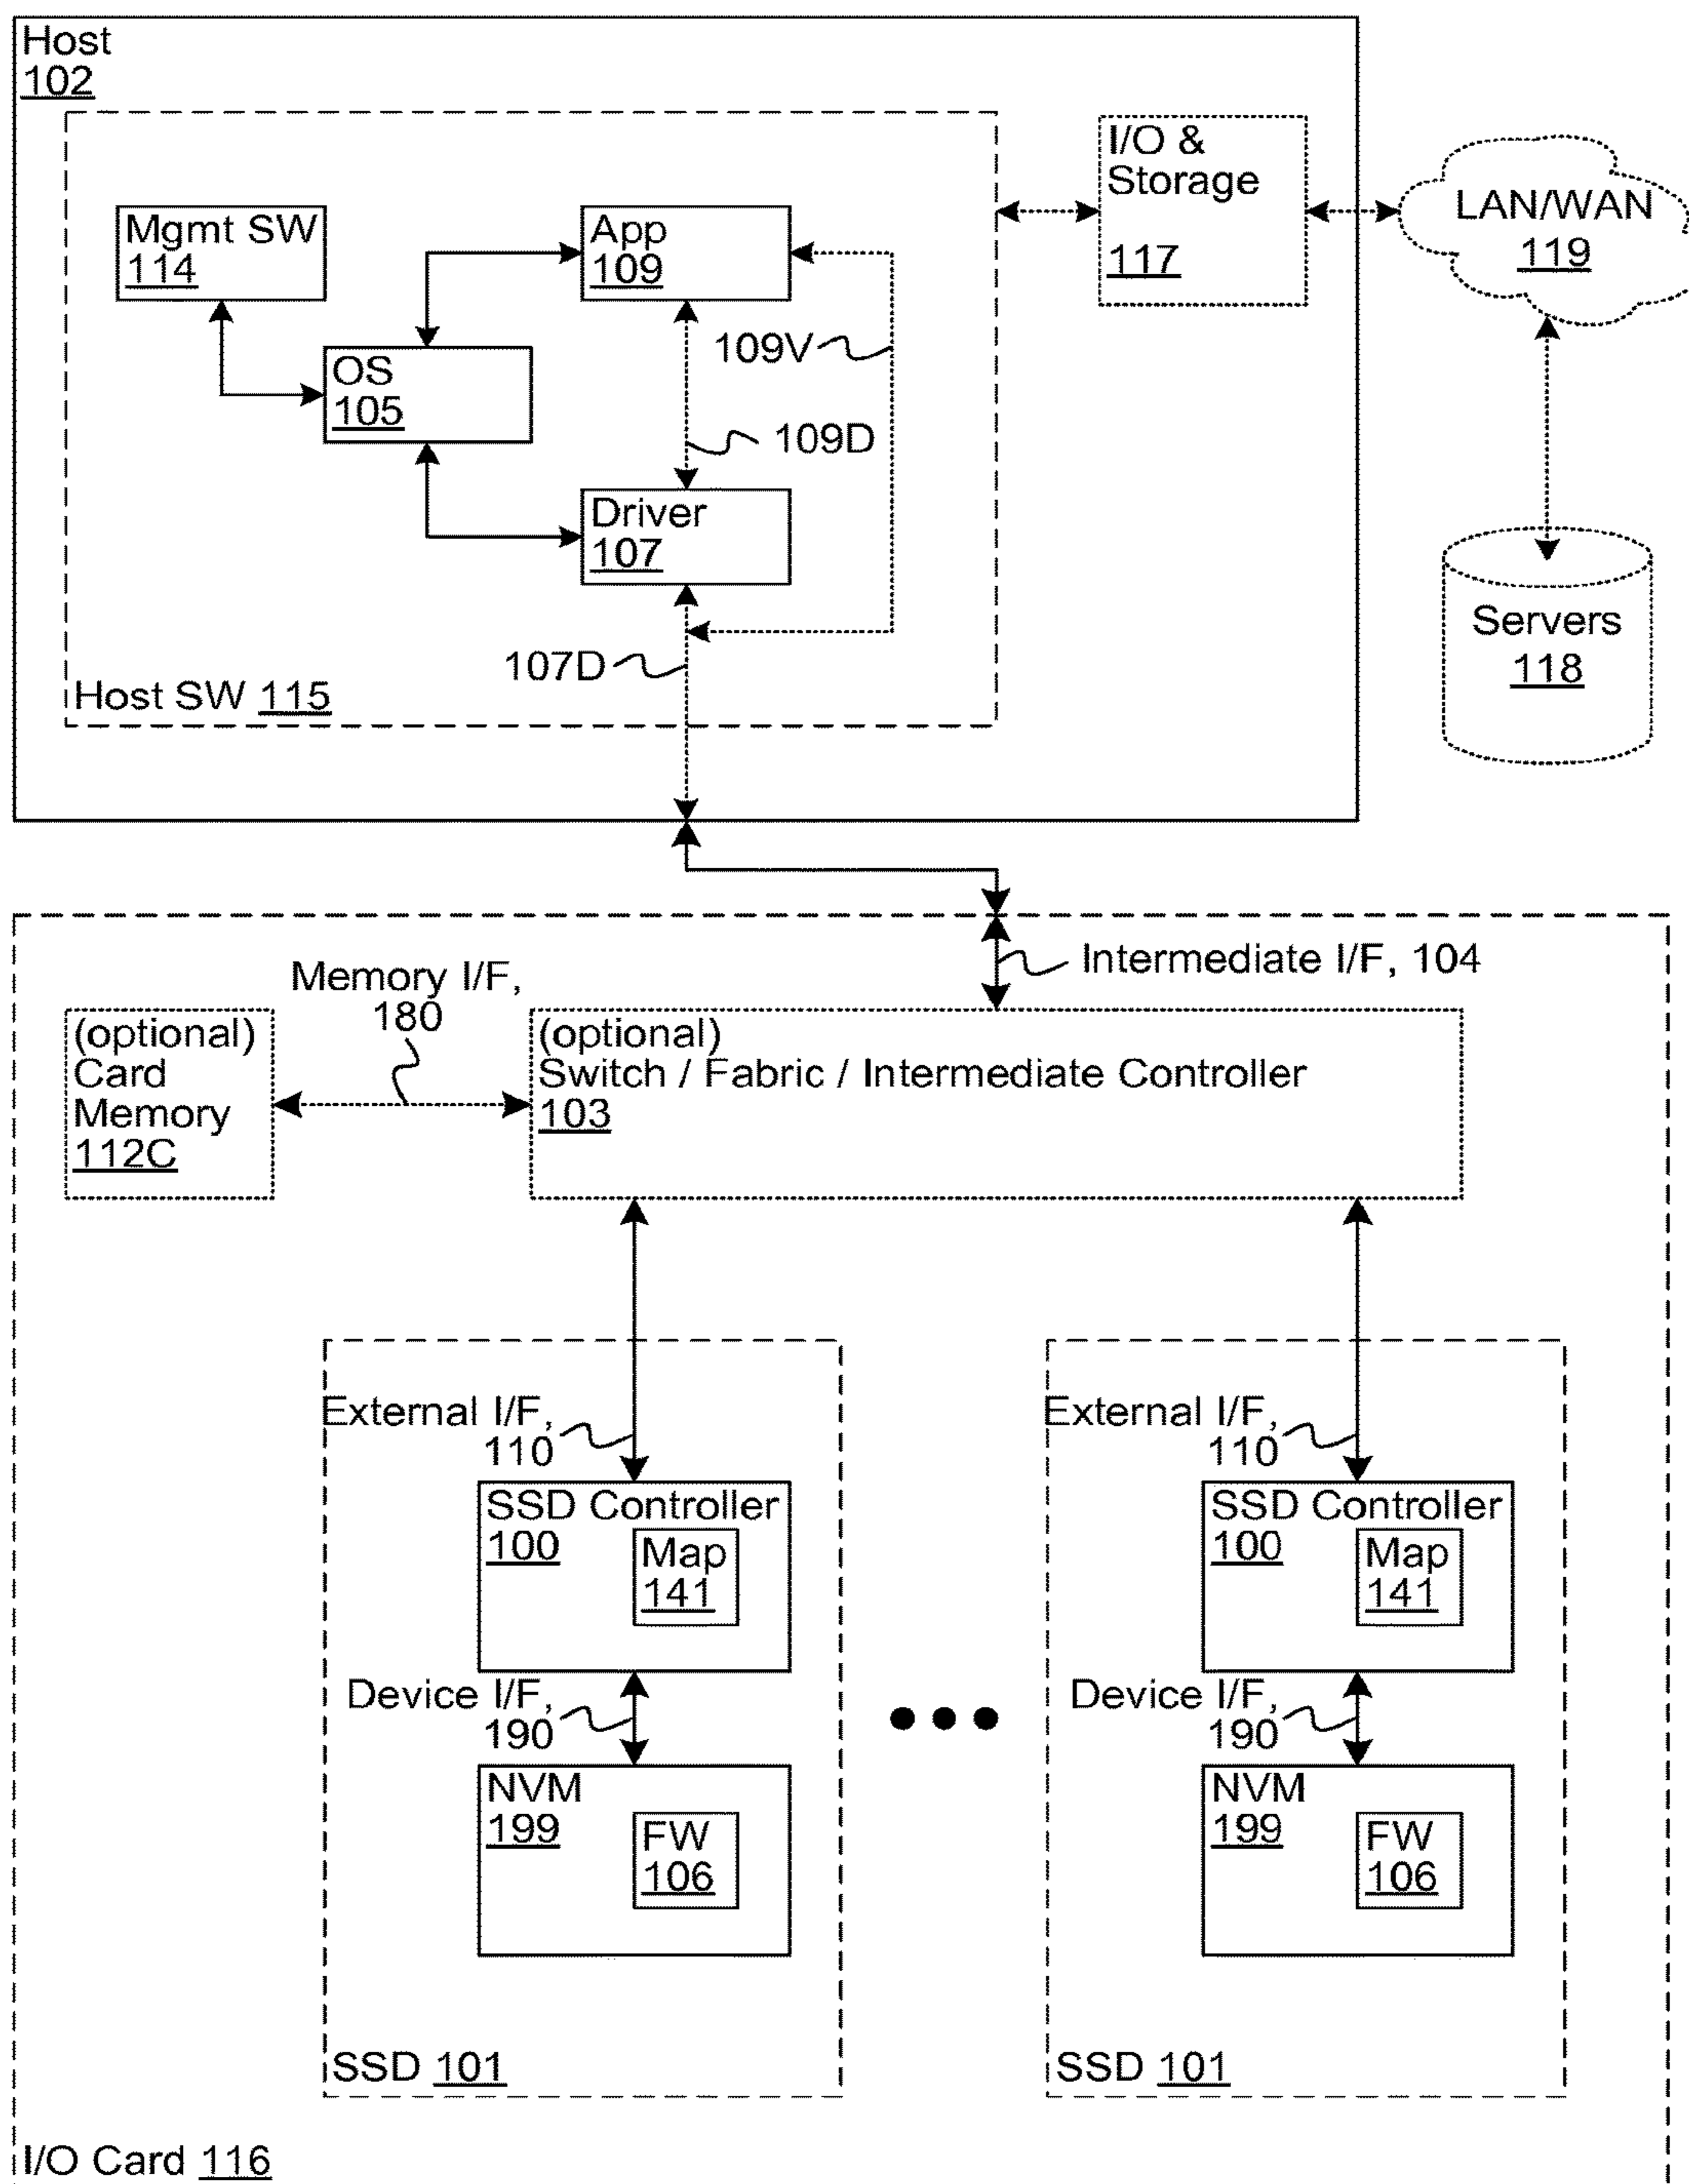


Fig. 1B

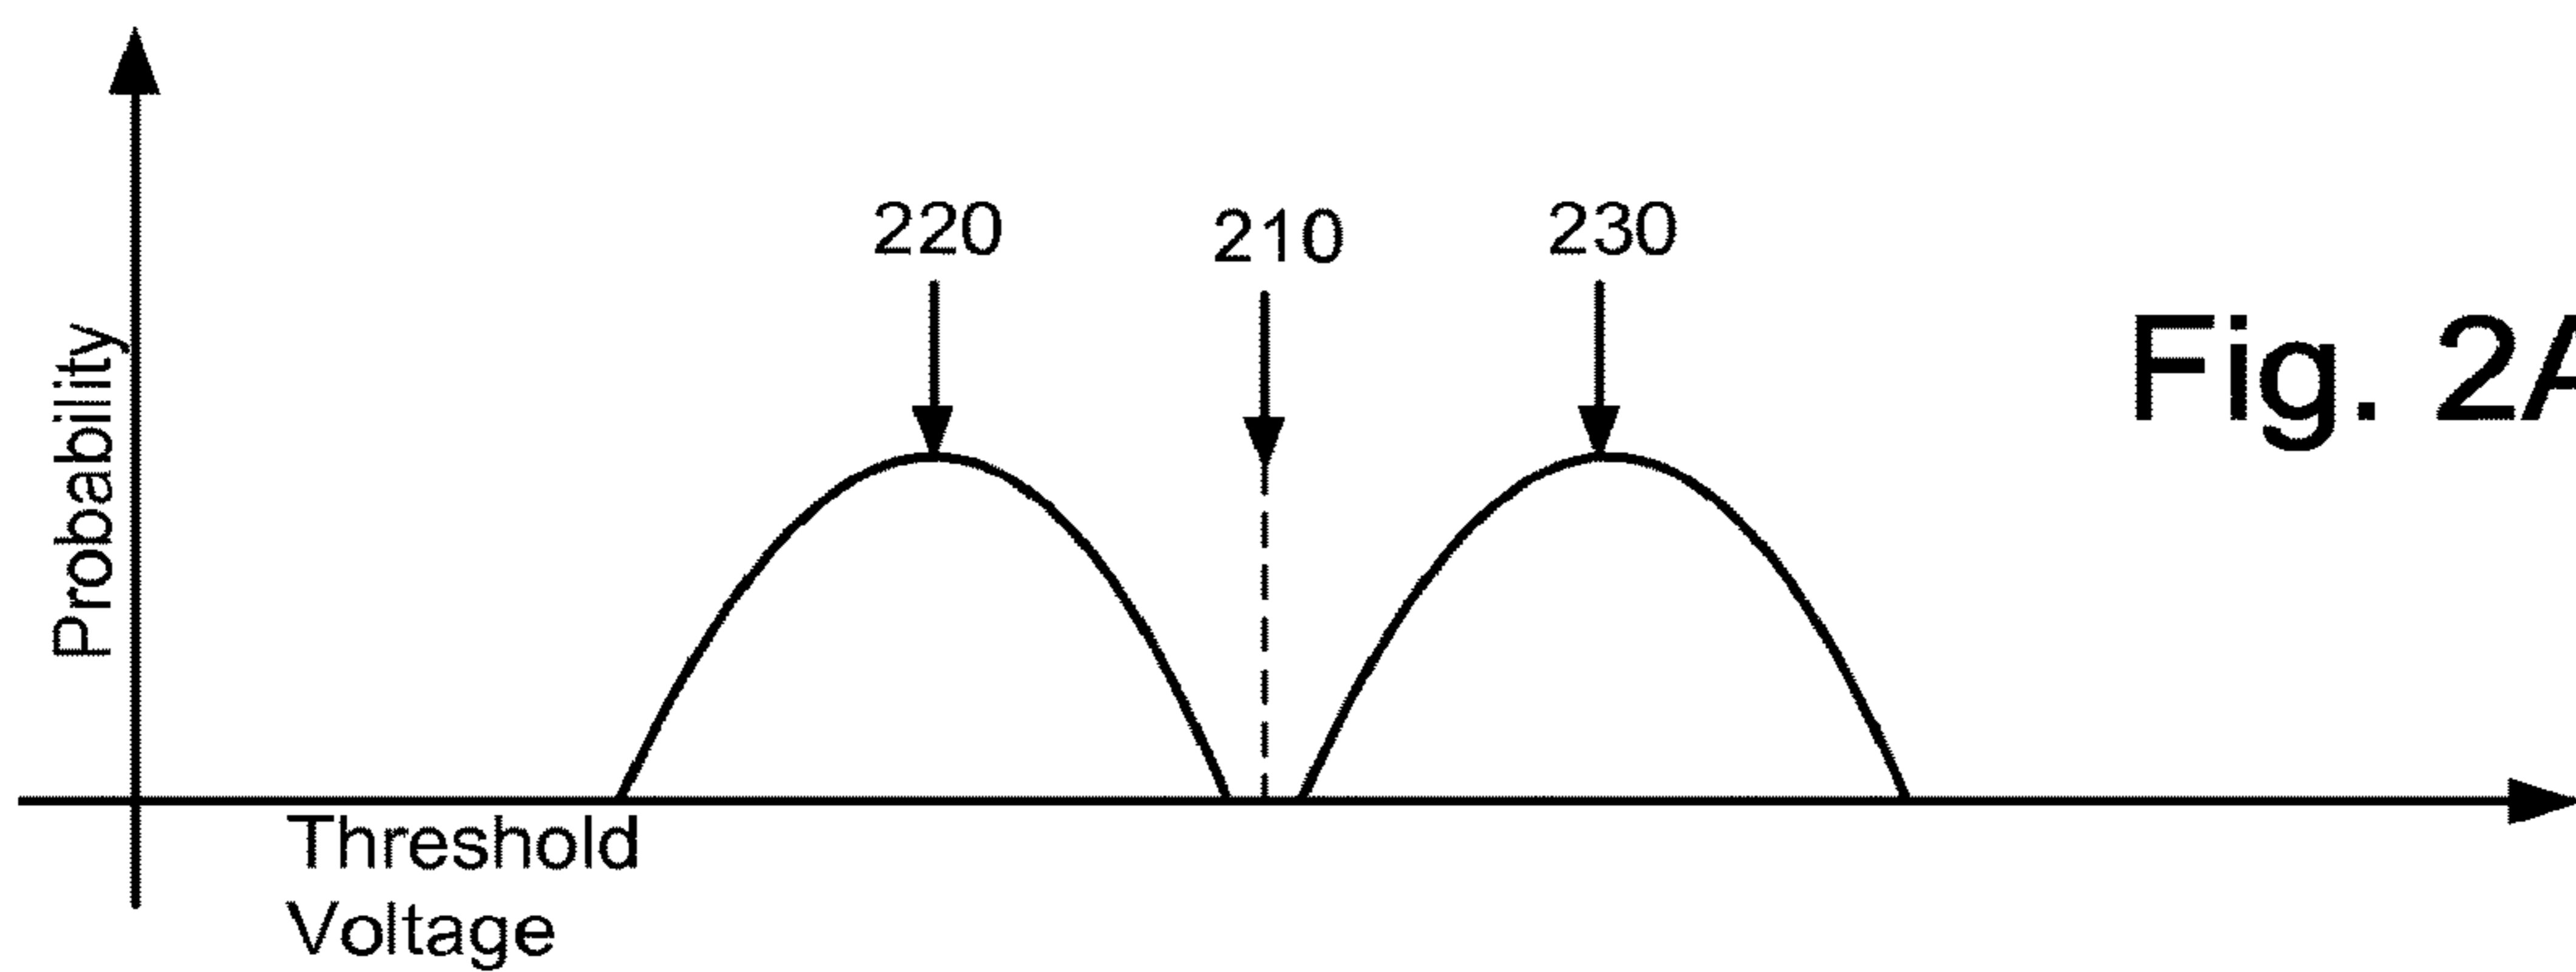


Fig. 2A

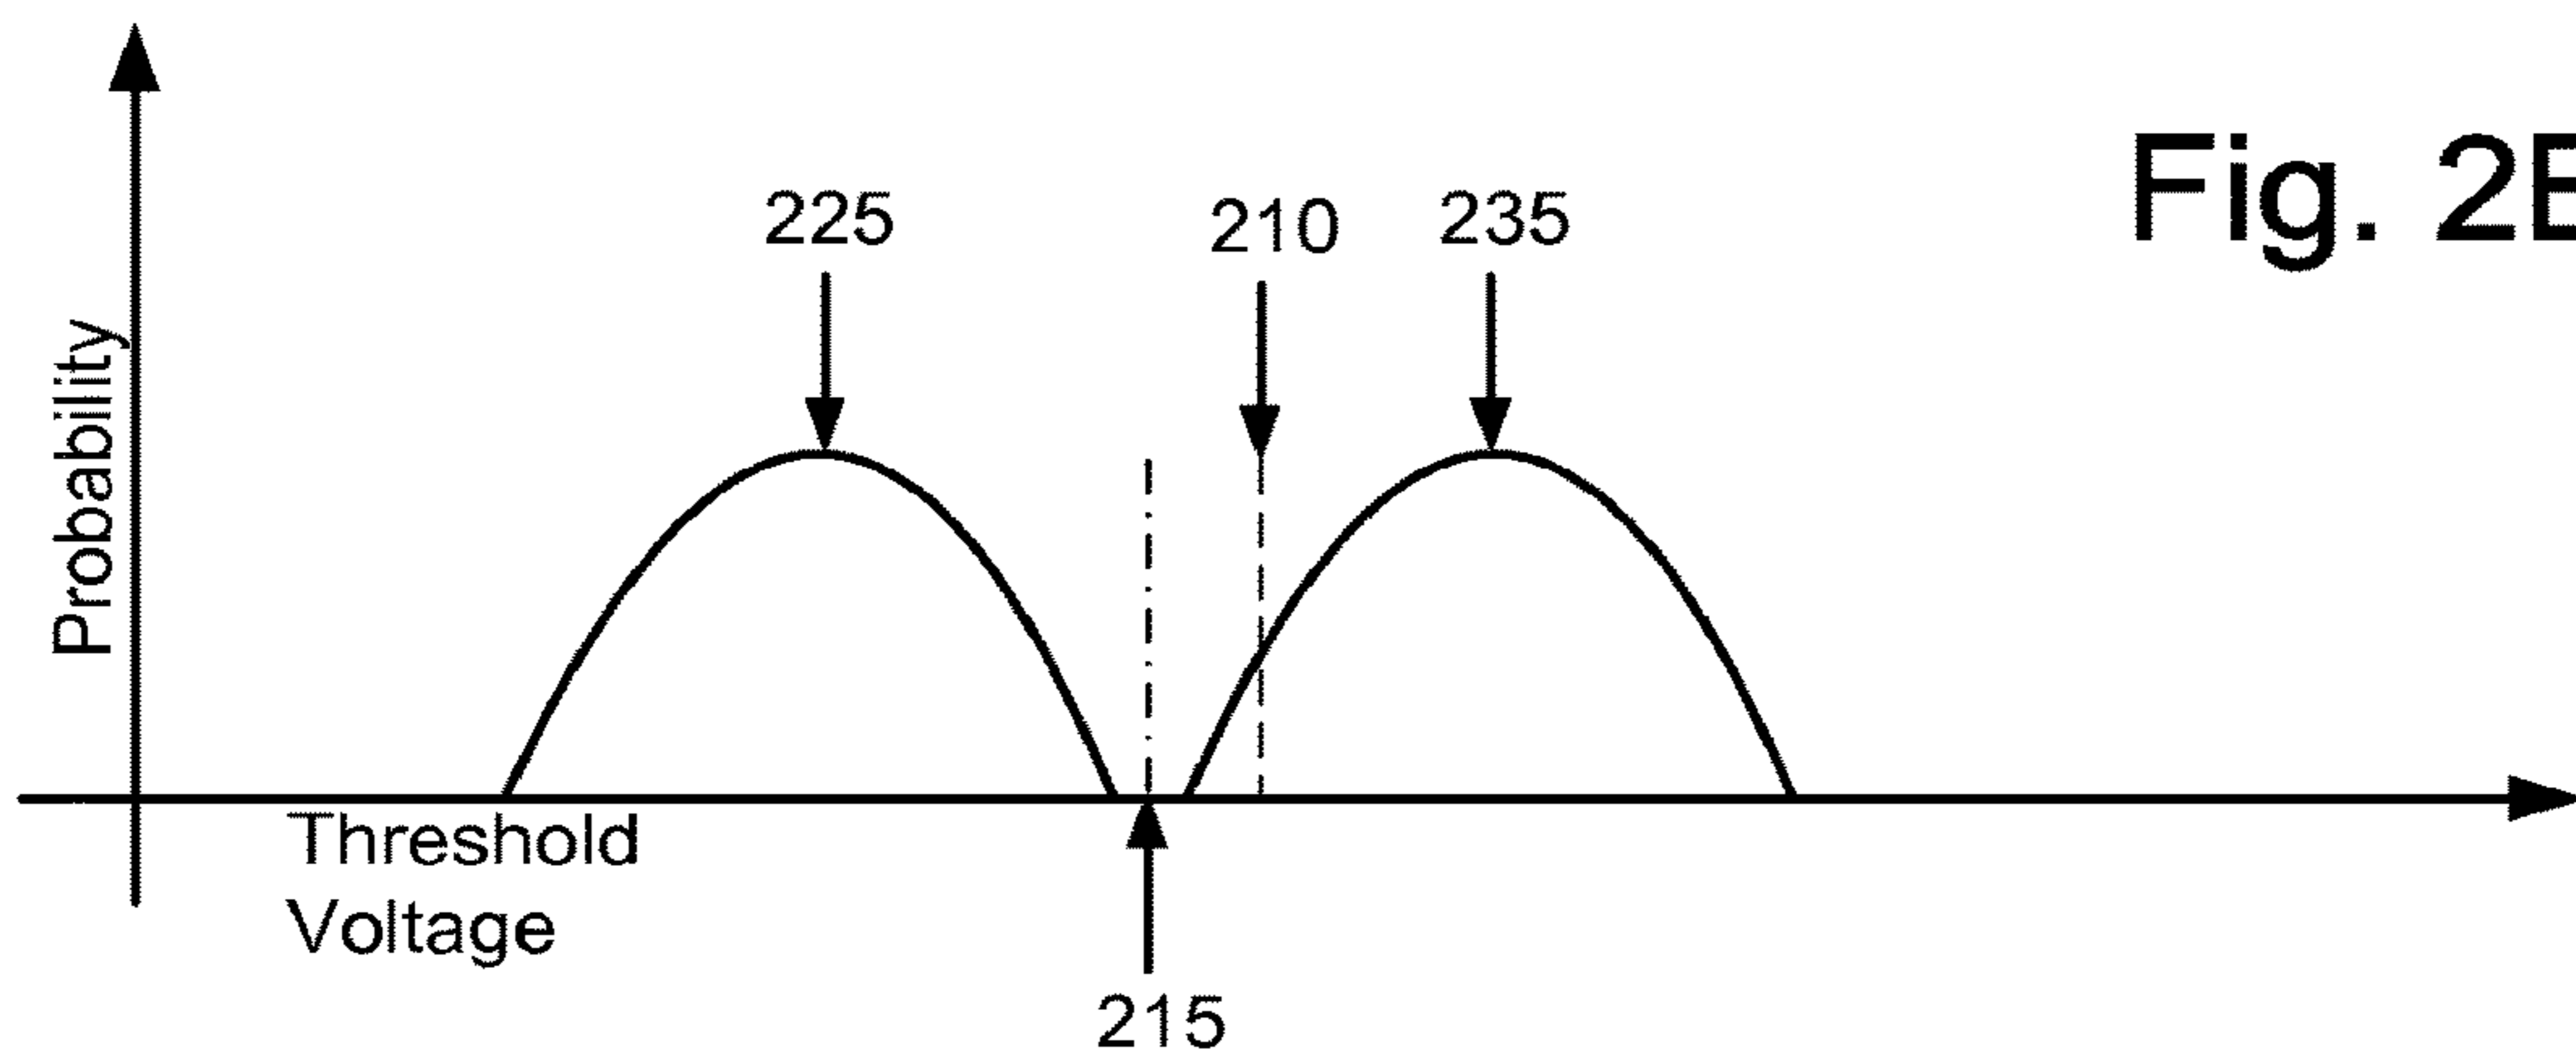


Fig. 2B

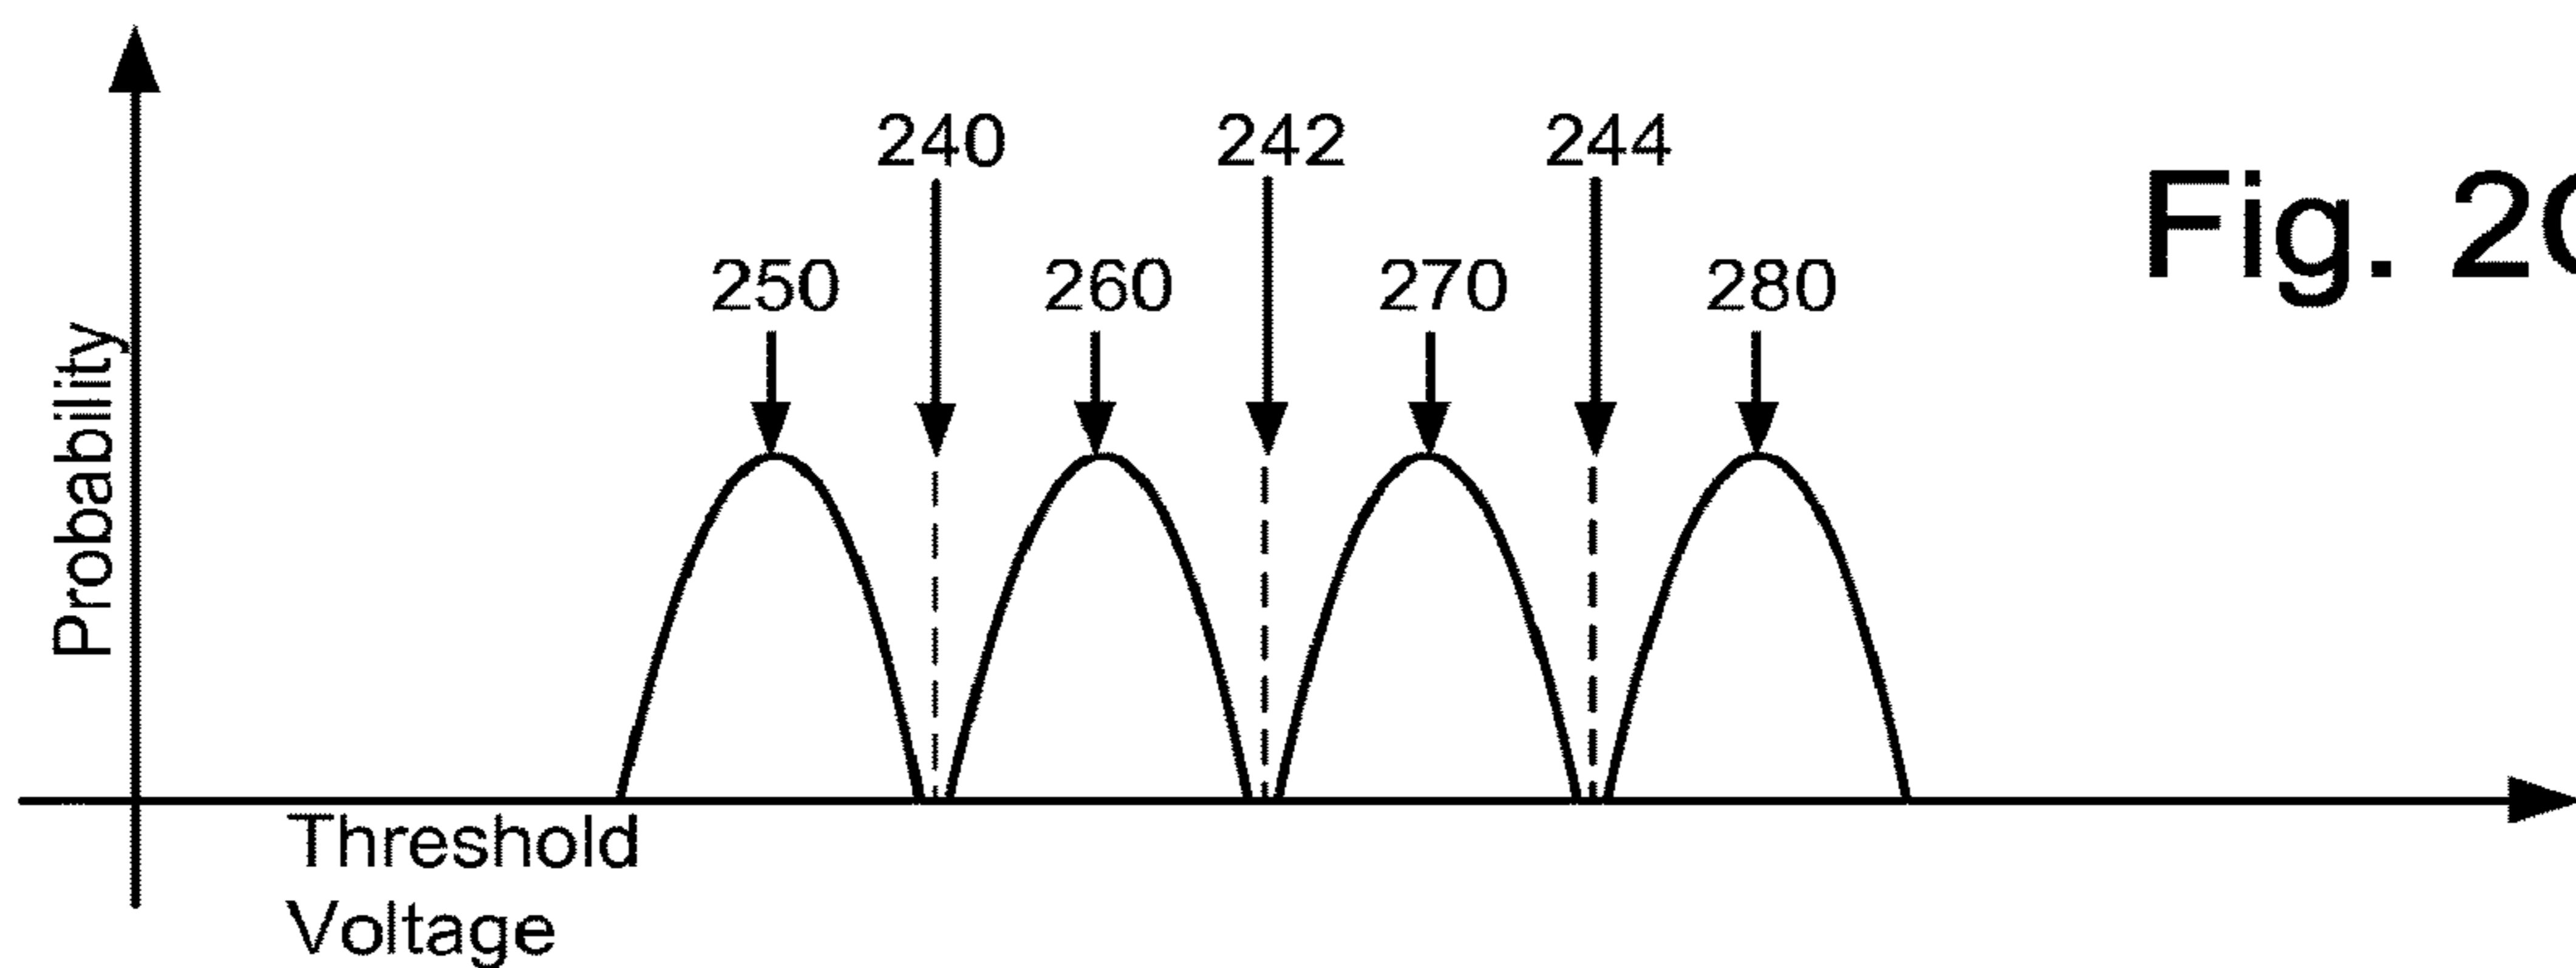


Fig. 2C

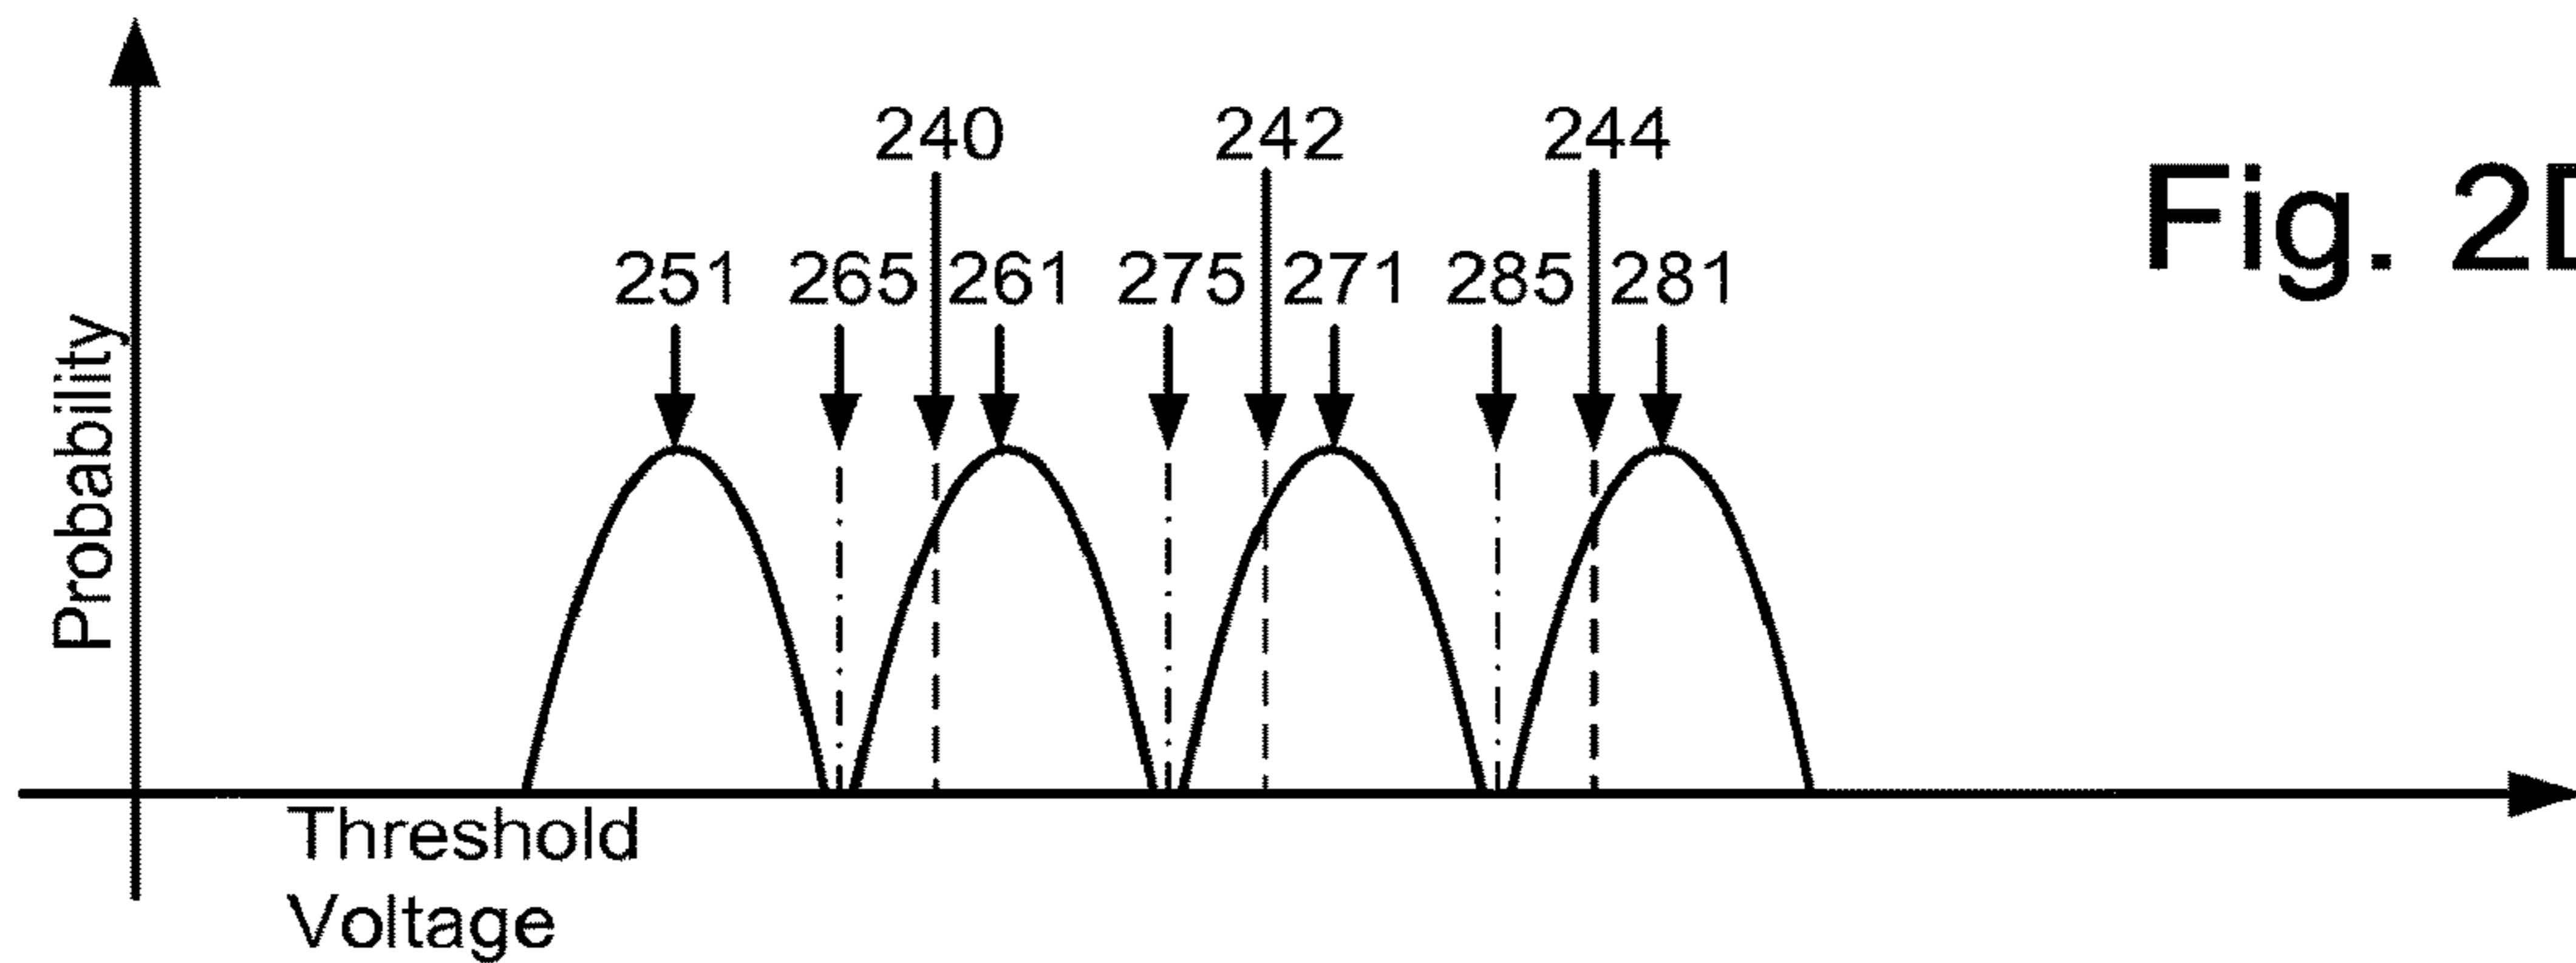


Fig. 2D

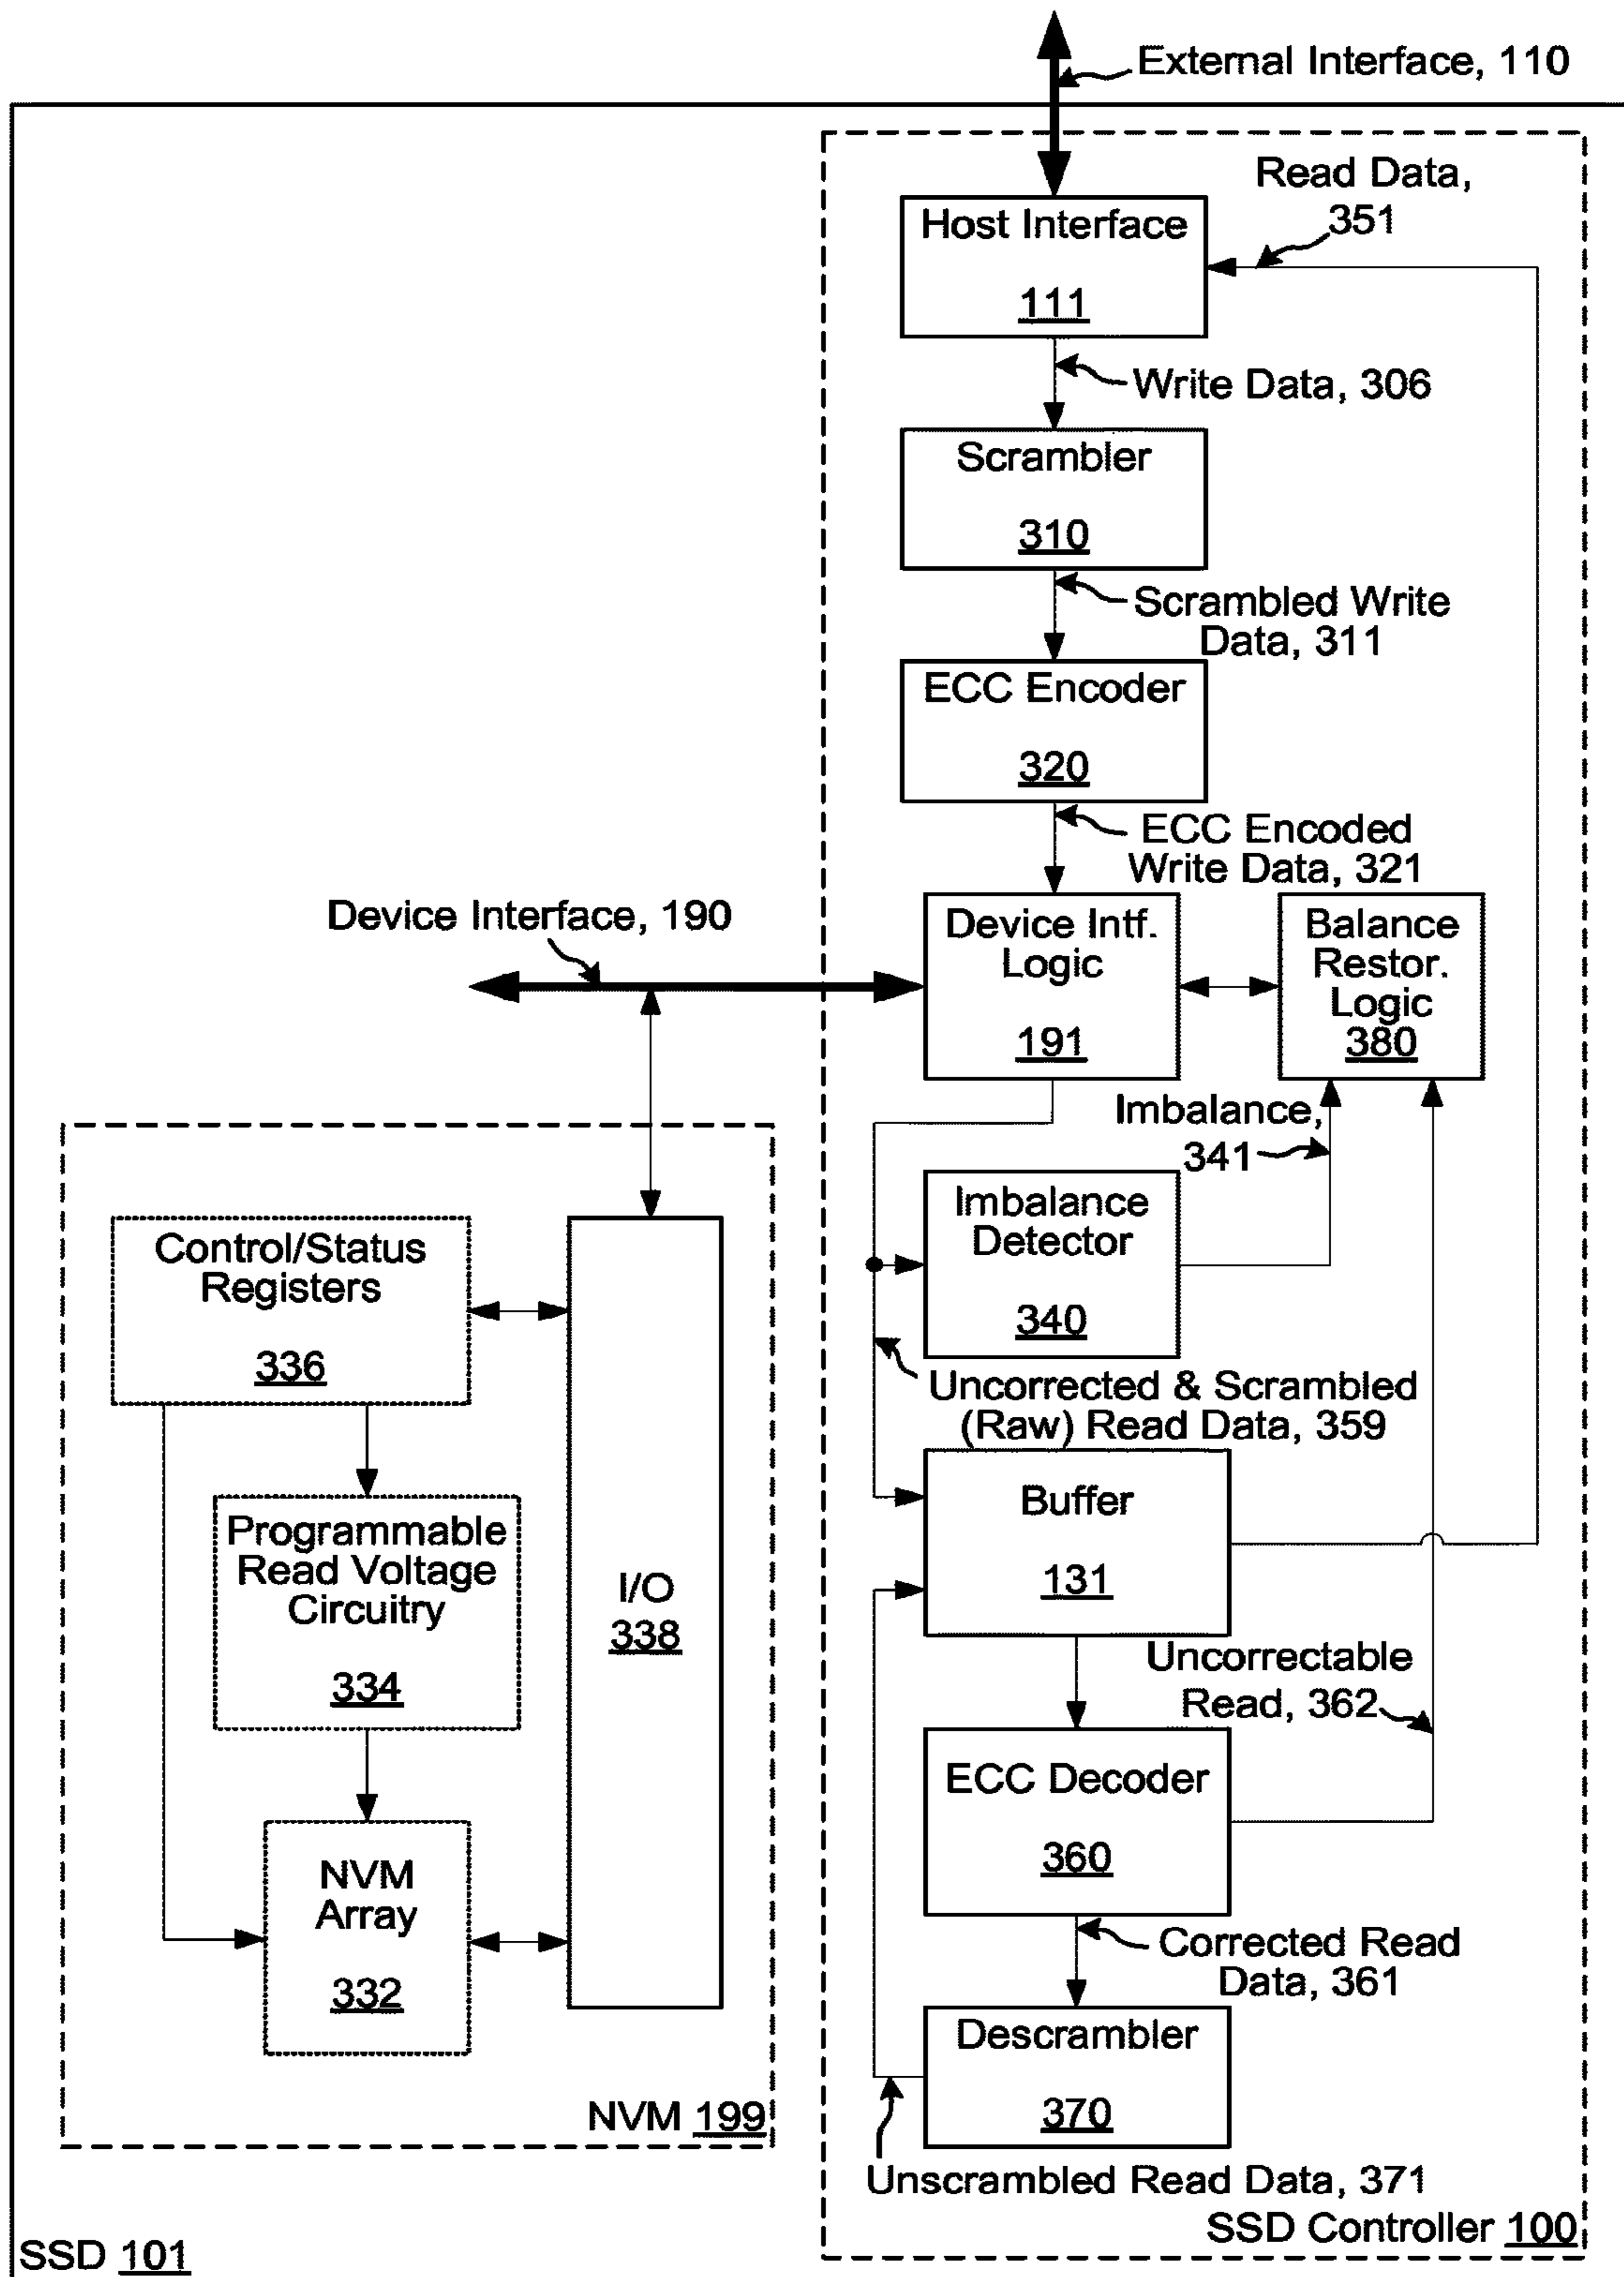


Fig. 3A

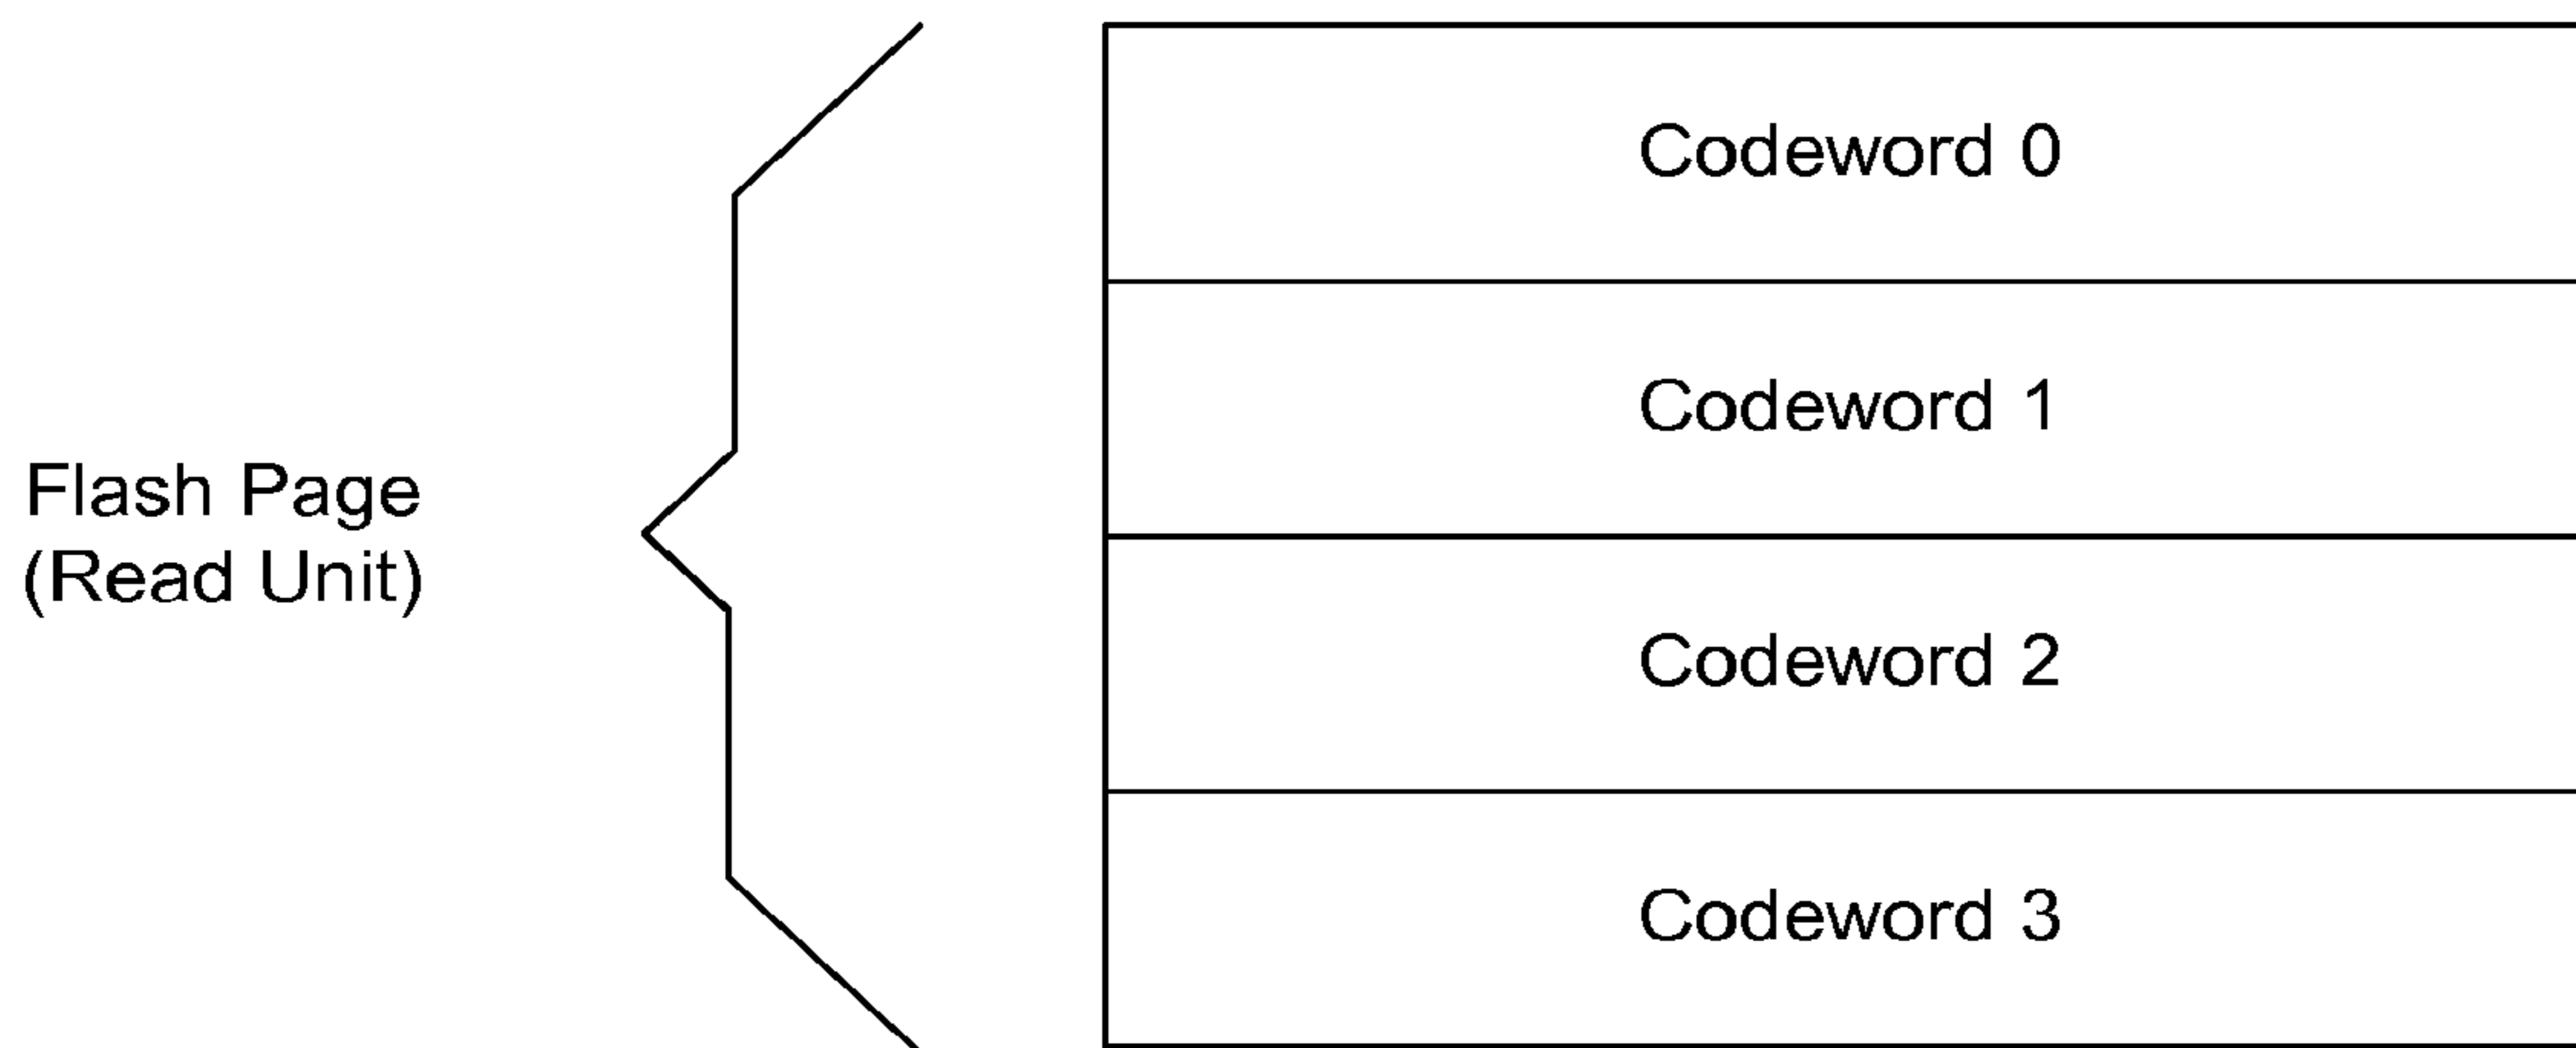


Fig. 3B

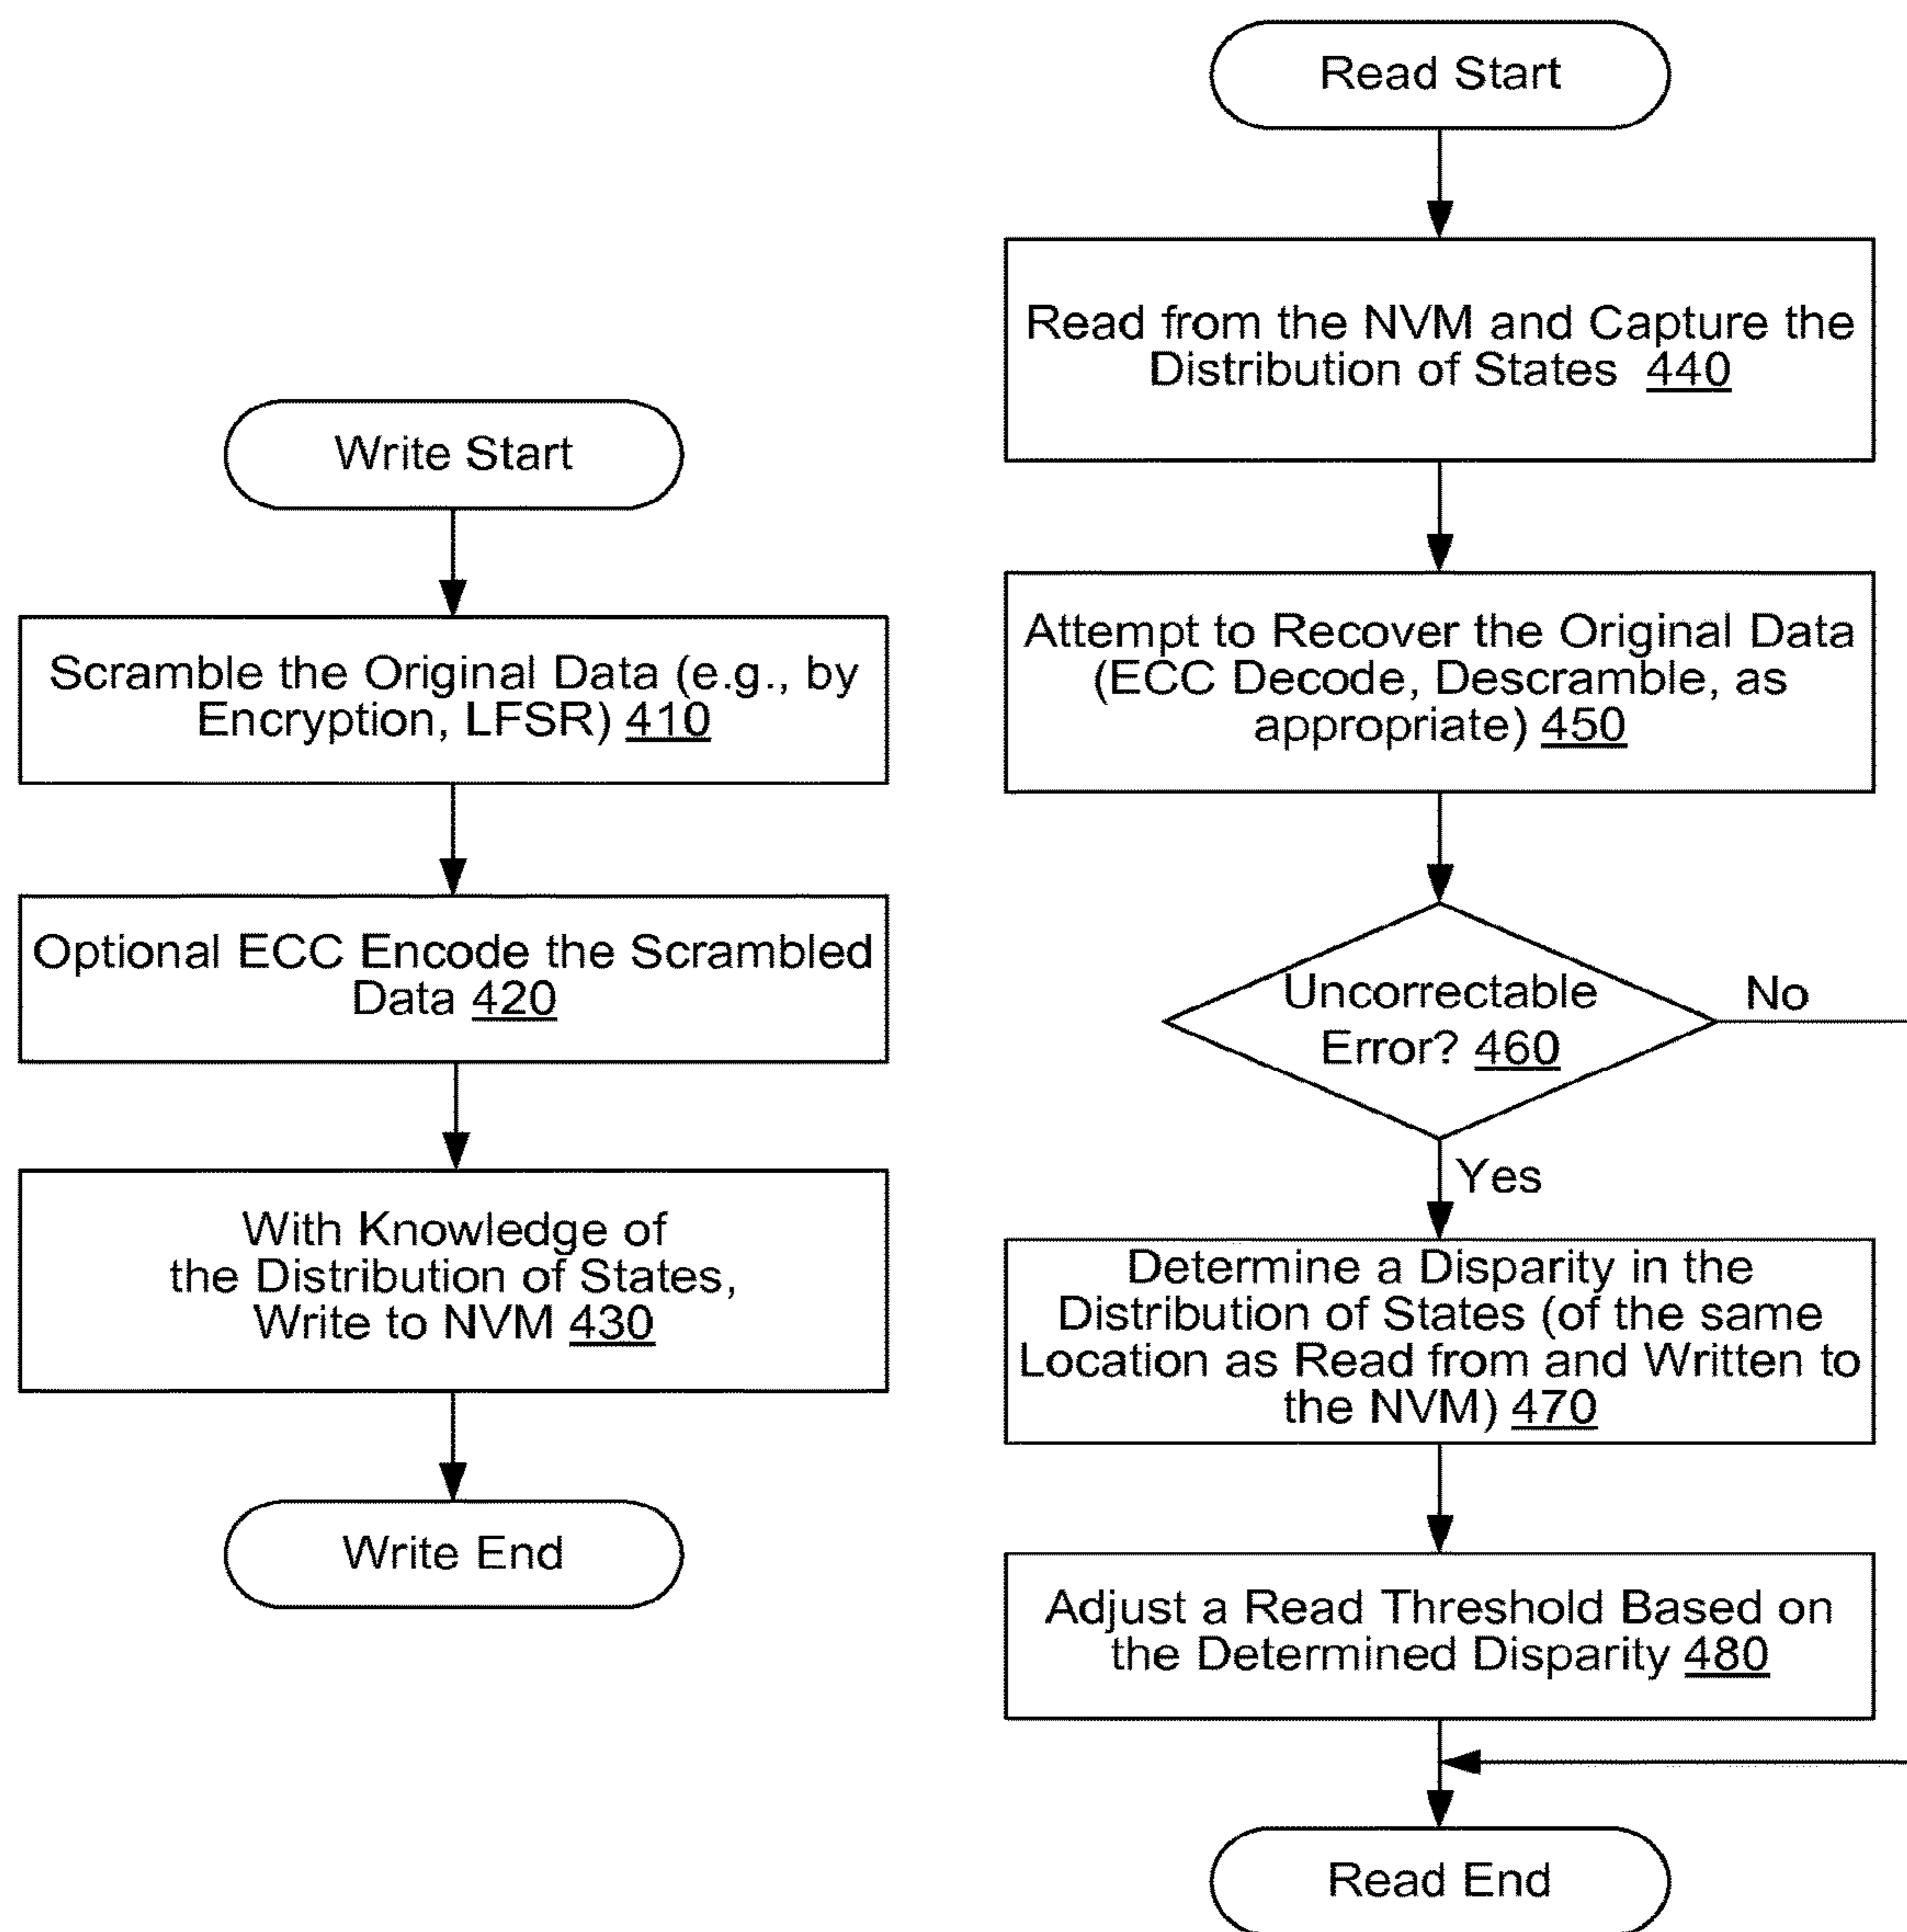


Fig. 4

ZERO-ONE BALANCE MANAGEMENT IN A SOLID-STATE DISK CONTROLLER

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 14/472,869 filed on Aug. 29, 2014, now U.S. Pat. No. 9,430,154, which is a continuation of U.S. patent application Ser. No. 13/464,433 filed on May 4, 2012, now U.S. Pat. No. 8,839,073, the entire disclosures of which are hereby specifically and entirely incorporated by reference.

BACKGROUND

Field

Advancements in non-volatile storage technology are needed to provide improvements in performance, efficiency, and utility of use.

Related Art

Unless expressly identified as being publicly or well known, mention herein of techniques and concepts, including for context, definitions, or comparison purposes, should not be construed as an admission that such techniques and concepts are previously publicly known or otherwise part of the prior art. All references cited herein (if any), including patents, patent applications, and publications, are hereby incorporated by reference in their entireties, whether specifically incorporated or not, for all purposes.

Various aspects of flash memory as used by Solid State Disk (SSD) controllers will now be described in part to establish a technology background and in part to establish antecedents for nomenclature used in the balance of the specification. The minimal size of data read by the SSD controller from Non-Volatile Memory (NVM) is the “read unit”, which is protected by included error correction, such as a Low-Density Parity-Check (LDPC) code. In some embodiments, each read unit contains approximately 4K to approximately 32K bits of user data, plus error correction overhead. Under command of the SSD controller, those bits are read from NVM memory cells, which depending on the technology as discussed below, may hold one or more bits per cell. In some embodiments, for security reasons an SSD controller encrypts the data prior to writing the data to NVM. In some embodiments, in view of circuit design limitations with respect to long strings of identically programmed cells, an SSD controller scrambles the data prior to writing the data to NVM.

Considered individually, each cell has a particular stored (programmed) charge that corresponds to a device threshold voltage for that cell, and further corresponds to the logical bit values being stored in the cell. While ideally all of the cells in the NVM would have identical device threshold voltages for the logical bit values stored, in practice for a variety of reasons the device threshold voltages differ across the cells in probability distributions along the threshold voltage axis (i.e., “threshold voltage distributions”) that are similar to a Gaussian in shape.

Thus considered in aggregate across a large number of cells, such as of a read unit, there are as many device threshold voltage distributions (e.g., Gaussian probability curves) as there are states per cell (two states per bit of storage per cell). That is, for N-bits per cell of storage, there are 2^N states and the same number of threshold voltage distributions. Generally, $(2^N)-1$ different read thresholds

(read voltage references V_{READ1} through $V_{READ(N-1)}$) are required by read circuits in the NVM to distinguish between the 2^N states.

Continuing from above, for Single-Level Cell (SLC) flash memories, $N=1$. SLC memories thus store one-bit per cell of storage, have two device threshold voltage distributions (one for zeroes and another for ones), and require a single read threshold, read voltage reference V_{READ1} . From lower to higher threshold voltages, the two threshold voltage distributions are known as the E (Erased) state and D1 (first Data) state. While arbitrary, a common mapping (coding) is to assign logical one to the E state and logical zero to the D1 state. Thus references to zeroes and ones are proxy references for respective decodings of the D1 state and the E state.

Continuing further from above, for Multi-Level Cell (MLC) flash memories, $N>1$. MLC memories thus store more than one bit per cell, have more than two device threshold voltage distributions, and require multiple different read thresholds to distinguish the distributions. For example, a 4LC memory stores 2 bits per cell, has four device threshold voltage distributions, and generally requires three read thresholds (read voltage references V_{READ1} , V_{READ2} , and V_{READ3}). From lower to higher threshold voltages, the four device threshold voltage distributions are known as the E (Erased), D1 (Data1), D2 (Data2), and D3 (Data3) states. While arbitrary, each of the four threshold voltage distributions is also mapped (addressed) in accordance with a particular binary sequence, such as a Gray code sequence. Thus references to one or more of the 11, 10, 00, and 01 states, are proxy references for respective decodings of the E, D1, D2, and D3 states.

With respect to address mapping used for the states of an MLC, each can be said to have a Most Significant Bit (MSB) and a Least Significant Bit (LSB) (and for more than 2 bits per cell, bits of significance in between). While there are various ways that MLC NVMs program their cells, the following approach is common. An initial programming round (a manipulation of the charge distributions) establishes the LSB, e.g., writes the “lower pages”. This is done loosely in the same manner as for writing an SLC, e.g., a charge manipulation that establishes the E state threshold voltage distribution and a second state threshold voltage distribution. Depending on the binary sequence used, the second state threshold voltage distribution is similar to the D1 state threshold voltage distribution, similar to the D2 state threshold voltage distribution, or between the D1 and D2 state threshold voltage distributions. For MLC, one or more additional programming rounds further manipulate the device threshold voltage distributions (in number, location along the voltage threshold axis, and in shape), as required per the number of levels of the MLC. More particularly, one or more subsequent programming operations write the “middle pages” (if any, for more than 2 bits per cell), and a last programming operation establishes the MSB, e.g., writes the “upper pages”. For example in a 4LC (2-bit per cell MLC), the E distribution and the second distribution of the first program round are respectively bifurcated by a second program round into E and D1 distributions and D2 and D3 distributions.

The threshold voltage distributions are modified away from their initial/nominal distributions by one or more factors, such as read disturb, write disturb, and retention loss. More particularly, over time, temperature, and other factors related to use, the location of each of the threshold voltage distribution can move around with respect to the threshold voltage axis. Such changes increase the likelihood

of read errors that are performed using a read reference voltage value for the read threshold that was previously established based on the nominal threshold voltage distribution. In some SLC embodiments, when a hard-decision uncorrectable error is encountered in a read unit read from NVM, a series of retry operations is performed to recover the read unit. The retry operations include the NVM re-reading (sampling) the read unit at different voltage values of the read threshold V_{READ1} , such as determined by a register setting written via an I/O command from the SSD controller. By reading at different settings of the read threshold V_{READ1} , the read unit is sampled at different points on the threshold voltage axis in an attempt to locate a sample of the read unit that is (hard-decision) correctable.

One algorithm suggested by NVM vendors is sweeping the SLC read threshold slowly upwards (increasing the read reference voltage V_{READ1} from its nominal value) in an attempt to find a sample of the read unit that is correctable. If this procedure fails, then the read threshold is swept in another direction (decreasing the read reference voltage V_{READ1} from its nominal value). If both sweeps fail, then the read unit is uncorrectable (by hard-decision decoding). Linearly sweeping the read threshold, which might have 16 to 64 steps at different respective read reference voltage settings, requires many time-consuming samples of the read unit (each with an attendant recovery time latency). Even when such a search is infrequently required, and thus not a major impact on average latencies, for applications with strict maximum latency requirements, including some database applications, such a time-consuming search is unacceptable.

SYNOPSIS

The invention may be implemented in numerous ways, including as a process, an article of manufacture, an apparatus, a system, a composition of matter, and a computer readable medium such as a computer readable storage medium (e.g., media in an optical and/or magnetic mass storage device such as a disk, or an integrated circuit having non-volatile storage such as flash storage) or a computer network wherein program instructions are sent over optical or electronic communication links. In this specification, these implementations, or any other form that the invention may take, may be referred to as techniques. The Detailed Description provides an exposition of one or more embodiments of the invention that enable improvements in performance, efficiency, and utility of use in the field identified above. The Detailed Description includes an Introduction to facilitate the more rapid understanding of the remainder of the Detailed Description. The Introduction includes Example Embodiments of one or more of systems, methods, articles of manufacture, and computer readable media in accordance with the concepts described herein. As is discussed in more detail in the Conclusions, the invention encompasses all possible modifications and variations within the scope of the issued claims.

The inventors believe they are the first to note the significance, with respect to read threshold management and attendant low latency performance and extended lifetime benefits, of the scrambling of data written to NVM and the tracking and managing of the statistical distribution of state of the scrambled data. More particularly, the inventors believe they are the first to note that a disparity in the statistical distributions of state of data read vs. data written is indicative of a shift in a preferred read threshold to use for reading the NVM. In some embodiments, an SSD controller

scrambles data, such as via encryption, prior to writing the data to NVM. The scrambled data has a known zero/one balance (statistical distribution of bits stored as zeroes and ones), such as having close to a 50-50 balance of zero bits and one bits. With efficient scrambling, the data is effectively random. In a random sample the size of an NVM read unit, a distribution of zeroes and ones tends to be very tight (narrowly dispersed about the mean). For example, with 16K (16,384) random bits, a probability of either the zeroes count or the ones count being more than approximately 300 away from a mean (8K) in either direction is <1% (and a probability of being more than approximately 500 away in either direction is <1 in a million).

Shifts in the various threshold voltage distributions with respect to the threshold voltage axis (due to any number of factors), away from a nominal distribution, act to change the location of a center-point between the threshold voltage distribution peaks. Changes in this actual “threshold voltage center-point” without corresponding tracking adjustments to the corresponding read threshold are a cause for read errors. The inventors believe they are the first to note the significance of this change of center-point, with respect to read threshold voltage management and attendant low latency performance and extended lifetime benefits.

Exploiting the above observations, in some embodiments an SSD controller at least temporarily maintains a zero count and a one count for each read unit read from an SLC NVM (or the lower pages of an MLC). In the event that the read unit is uncorrectable in part due to the shift in the threshold voltage distributions away from their (initial) nominal distributions, the zero count and/or the one count enable a determination of a direction and/or a magnitude to move a read threshold (read voltage reference V_{READ1}) to track the shifted threshold voltage distributions and restore the zero/one balance. As presented in the Detailed Description, according to embodiment, the new voltage setting for the read threshold is determined in a variety of ways (counts, percentages) based on a number of factors (e.g. observed/inferred threshold voltage distributions, known stored values, and past NVM operating events).

The above techniques are extended for MLC memories, including maintaining counts for upper pages and shifting of multiple read thresholds. The above techniques are also usable with a variety of encryption/scrambling methods having known zero/one balances other than a 50-50 balance.

BRIEF DESCRIPTION OF DRAWINGS

FIG. 1A illustrates selected details of an embodiment of a Solid-State Disk (SSD) including an SSD controller providing a platform to implement zero/one balance management techniques in non-volatile storage, such as implemented via Non-Volatile Memory (NVM) elements (e.g., flash memories).

FIG. 1B illustrates selected details of various embodiments of systems including one or more instances of the SSD of FIG. 1A.

FIGS. 2A through 2D illustrate example threshold voltage distributions relevant to zero/one balance restoration. FIGS. 2A and 2B are SLC specific, respectively illustrating threshold voltage distributions at a first and a later second time. FIGS. 2C and 2D are MLC specific, respectively illustrating threshold voltage distributions at a first and a later second time.

FIG. 3A illustrates selected details of a system embodiment of an instance of the SSD of FIG. 1A, providing particular details regarding zero/one balance management.

5

FIG. 3B illustrates selected details of the organization of a read unit used by the embodiment of FIG. 3A.

FIG. 4 illustrates selected control flow details for an embodiment of the SSD of FIGS. 1A, 1B, and 3, providing particular details regarding write related operations (actions 410 through 430) and read related operations (actions 440 through 480).

List of Reference Symbols in Drawings	
Ref. Symbol	Element Name
100	SSD Controller
101	SSD
102	Host
103	(optional) Switch/Fabric/Intermediate Controller
104	Intermediate Interfaces
105	OS
106	FirmWare (FW)
107	Driver
107D	dotted-arrow (Host Software $\leftarrow \rightarrow$ I/O Device Communication)
109	Application
109D	dotted-arrow (Application $\leftarrow \rightarrow$ I/O Device Communication via driver)
109V	dotted-arrow (Application $\leftarrow \rightarrow$ I/O Device Communication via VF)
110	External Interfaces
111	Host Interfaces
112C	(optional) Card Memory
113	Tag Tracking
114	Multi-Device Management Software
115	Host Software
116	I/O Card
117	I/O & Storage Devices/Resources
118	Servers
119	LAN/WAN
121	Data Processing
123	Engines
131	Buffer
133	DMA
135	ECC-X
137	Memory
141	Map
143	Table
151	Recycler
161	ECC
171	CPU
172	CPU Core
173	Command Management
175	Buffer Management
177	Translation Management
179	Coherency Management
180	Memory Interface
181	Device Management
182	Identity Management
190	Device Interfaces
191	Device Interface Logic
192	Flash Device
193	Scheduling
194	Flash Die
199	NVM
210	SLC Nominal Read Threshold
215	SLC Adjusted Read Threshold
220	SLC Nominal Most-Negative Threshold State
225	SLC Shifted Most-Negative Threshold State
230	SLC Nominal Most-Positive Threshold State
235	SLC Shifted Most-Positive Threshold State
240	MLC Nominal First Read Threshold
242	MLC Nominal Second Read Threshold
244	MLC Nominal Third Read Threshold
250	MLC Nominal Most-Negative Threshold State
251	MLC Adjusted Most-Negative Threshold State
260	MLC Nominal Next-Most-Negative Threshold State
261	MLC Adjusted Next-Most-Negative Threshold State
265	MLC Adjusted First Read Threshold
270	MLC Nominal Next-Most-Positive Threshold State
271	MLC Adjusted Next-Most-Positive Threshold State
275	MLC Adjusted Second Read Threshold

6

-continued

List of Reference Symbols in Drawings	
Ref. Symbol	Element Name
280	MLC Nominal Most-Positive Threshold State
281	MLC Adjusted Most-Positive Threshold State
285	MLC Adjusted Third Read Threshold
306	Write Data
310	Scrambler
311	Scrambled Write Data
320	ECC Encoder
321	ECC Encoded Write Data
332	NVM Array
334	Programmable Read Voltage Circuitry
336	Control/Status Registers
338	I/O
340	Imbalance Detector
341	Imbalance
351	Read Data
359	Uncorrected & Scrambled (Raw) Read Data
360	ECC Decoder
361	Corrected Read Data
362	Uncorrectable Read
370	Descrambler
371	Unscrambled Read Data
380	Balance Restoration Logic
410	Scramble Data
420	ECC Encode
430	Write NVM
440	Read NVM
450	Data Recovery
460	Uncorrectable Error Evaluation
470	Disparity Determination
480	Read Threshold Adjustment

DETAILED DESCRIPTION

35 A detailed description of one or more embodiments of the invention is provided below along with accompanying figures illustrating selected details of the invention. The invention is described in connection with the embodiments. The

40 embodiments herein are understood to be merely exemplary, the invention is expressly not limited to or by any or all of the embodiments herein, and the invention encompasses numerous alternatives, modifications, and equivalents. To avoid monotony in the exposition, a variety of word labels

45 (including but not limited to: first, last, certain, various, further, other, particular, select, some, and notable) may be applied to separate sets of embodiments; as used herein such labels are expressly not meant to convey quality, or any form of preference or prejudice, but merely to conveniently

50 distinguish among the separate sets. The order of some operations of disclosed processes is alterable within the scope of the invention. Wherever multiple embodiments serve to describe variations in process, method, and/or

55 program instruction features, other embodiments are contemplated that in accordance with a predetermined or a dynamically determined criterion perform static and/or dynamic selection of one of a plurality of modes of operation corresponding respectively to a plurality of the multiple

60 embodiments. Numerous specific details are set forth in the following description to provide a thorough understanding of the invention. The details are provided for the purpose of example and the invention may be practiced according to the claims without some or all of the details. For the purpose of

65 clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the invention is not unnecessarily obscured.

Introduction

This introduction is included only to facilitate the more rapid understanding of the Detailed Description; the invention is not limited to the concepts presented in the introduction (including explicit examples, if any), as the paragraphs of any introduction are necessarily an abridged view of the entire subject and are not meant to be an exhaustive or restrictive description. For example, the introduction that follows provides overview information limited by space and organization to only certain embodiments. There are many other embodiments, including those to which claims will ultimately be drawn, discussed throughout the balance of the specification.

Acronyms

At least some of the various shorthand abbreviations (e.g. acronyms) defined here refer to certain elements used herein.

Acronym	Description
AES	Advanced Encryption Standard
AHCI	Advanced Host Controller Interface
API	Application Program Interface
ASCII	American Standard Code for Information Interchange
ATA	Advanced Technology Attachment (AT Attachment)
BCH	Bose Chaudhuri Hocquenghem
CD	Compact Disk
CF	Compact Flash
CMOS	Complementary Metal Oxide Semiconductor
CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
DAS	Direct Attached Storage
DDR	Double-Data-Rate
DES	Data Encryption Standard
DMA	Direct Memory Access
DNA	Direct NAND Access
DRAM	Dynamic Random Access Memory
DVD	Digital Versatile/Video Disk
DVR	Digital Video Recorder
ECC	Error-Correcting Code
eMMC	Embedded MultiMediaCard
eSATA	external Serial Advanced Technology Attachment
GPS	Global Positioning System
HDD	Hard Disk Drive
I/O	Input/Output
IC	Integrated Circuit
IDE	Integrated Drive Electronics
JPEG	Joint Photographic Experts Group
LAN	Local Area Network
LB	Logical Block
LBA	Logical Block Address
LDPC	Low-Density Parity-Check
LFSR	Linear Feedback Shift Register
LPN	Logical Page Number
LSB	Least Significant Bit
LZ	Lempel-Ziv
MLC	Multi-Level Cell
MMC	MultiMediaCard
MPEG	Moving Picture Experts Group
MSB	Most Significant Bit
NAS	Network Attached Storage
NCQ	Native Command Queuing
NVM	Non-Volatile Memory
ONA	Optimized NAND Access
ONFI	Open NAND Flash Interface
OS	Operating System
PC	Personal Computer
PCIe	Peripheral Component Interconnect express (PCI express)
PDA	Personal Digital Assistant
PHY	PHYSical interface
POS	Point Of Sale
RAID	Redundant Array of Inexpensive/Independent Disks
RASIE	Redundant Array of Silicon Independent Elements
ReRAM	Resistive Random Access Memory
RS	Reed-Solomon
RSA	Rivest, Shamir & Adleman
SAN	Storage Attached Network
SAS	Serial Attached Small Computer System Interface

-continued

Acronym	Description
	(Serial SCSI)
5 SATA	Serial Advanced Technology Attachment (Serial ATA)
SCSI	Small Computer System Interface
SD	Secure Digital
SDR	Single-Data-Rate
SLC	Single-Level Cell
SMART	Self-Monitoring Analysis and Reporting Technology
10 SSD	Solid-State Disk/Drive
TCG	Trusted Computing Group
UFS	Unified Flash Storage
USB	Universal Serial Bus
VF	Virtual Function
15 WAN	Wide Area Network

Tracking and Restoring Zero/One Balances

The initial discussion that follows, of tracking and managing counts for zeroes and ones and the shifting of the read threshold used to distinguish the zero and one states, is directly applicable to SLC memories and (with minor variation, as noted) to the lower pages of MLC memories. More generally with respect to MLC memories, extension to tracking and managing counts for upper pages and the shifting of multiple read thresholds (e.g., read voltage references V_{READ1} , V_{READ2} , and V_{READ3}) is detailed in separate paragraphs below.

In some embodiments, a zero count and a one count are obtained for each read unit (or each group of read units transferred at one time from one NVM die) read from an SLC NVM. In an event that the read unit is uncorrectable in part due to the shift in the threshold voltage distributions away from their (initial) nominal distributions, the zero count and/or the one count enable a determination of a direction and/or a magnitude to move (set) a read threshold to track the shift in the threshold voltage distributions and restore the zero/one balance. The read threshold being adjusted is read voltage reference V_{READ1} for SLC memories and read voltage reference V_{READ2} for the lower pages of MLC memories. For example, in an SLC having a logical one assigned to the erase state, and written with a known statistically equal (50-50 percent, a.k.a. 50/50 percent, or simply 50/50) distribution of zeroes and ones, if a current zero/one counts indicates a predominance of ones, V_{READ1} is shifted lower; and if the current zero/one counts indicates a predominance of zeroes, V_{READ1} is shifted higher. (Substitute V_{READ2} for V_{READ1} in the foregoing example, for the lower pages of MLC memories.) A magnitude of how much to change (shift) the read threshold is optionally determined by a percentage of a number of zeroes (or ones). The magnitude is based on one or more of: the number of zeroes and/or ones observed in one or more samples; a read threshold voltage increment of the one or more samples; a knowledge of predicted threshold voltage distributions corresponding to a midpoint between the peaks of the two distributions; program/erase count (wear); read disturb counts; retention time; and other similar factors.

Finding the Read-Equilibrium Point

In further embodiments, a binary-search-like operation (iteratively sampling the same read unit at respective values of the appropriate read voltage reference) is used to find a read threshold "read-equilibrium point". This is a point on the threshold voltage axis between adjacent threshold voltage distributions resulting in raw data (pre error correction) statistical state distributions that upon reading, match within 65 a statistical margin of error to the statistical state distributions written, or idealized models of same.

For SLC memory and random scrambling, such as resulting from encryption, the read-equilibrium point is the read threshold (read voltage reference V_{READ1}) where a zero/one balance (of the raw data read from the NVM) is most even (near identical). Under such conditions, the two state distributions have a 50-50 balance. I.e., 50% of the states read are zeroes and 50% of the states read are ones. (Substitute V_{READ2} for V_{READ1} in the foregoing, for the lower pages of MLC memories.)

In some embodiments, the read-equilibrium point corresponds to one or more of: a center point between two voltage distributions, a minima between two voltage distributions, a point at which the zero/one balance in read data is closest to 50-50, and a point determined according to interpolation of the zero/one balance found at two other points. The read equilibrium point corresponds to a center-point between the adjacent threshold voltage distributions in embodiments having symmetric adjacent threshold voltage distributions.

Finding and setting the read threshold to the desired read-equilibrium point reduces a number of reads required to find an optimal point at which to attempt hard-decision decoding. In still further embodiments, hard-decision decoding is attempted at each read threshold sampling point (each value of V_{READ1} for SLC memories, or each value of V_{READ2} for the lower pages of MLC memories) in the binary search in the event a “good enough” read threshold is found before the read threshold read-equilibrium point is found. The search is generally truncated upon finding the good enough read threshold, except as discussed below regarding possibly taking additional sampling around the determined read-equilibrium point for soft-decision decoding purposes. In various embodiments, the “good enough” read threshold results in successful hard decision decoding of the read data. In some embodiments, a magnitude of the search steps (the read threshold voltage increment) in the binary search is determined (based), at least in part, on various factors as detailed above regarding the magnitude of the read threshold voltage change.

MLC Considerations

In MLC memories, multiple read thresholds are managed. In some embodiments this is performed by assuming uniform displacements of the multiple device-threshold-voltage distributions and changing all other read thresholds based on decisions made for the first read threshold (V_{READ3}) based on a read data zero/one balance of the lower pages.

For a 4LC NVM, theoretically there is an 11/10/00/01 balance of 25-25-25-25 (in contrast to a zero/one balance of 50-50). However, there is no single operation that directly provides this balance, given that 4LCs are generally read via two separate read operations using the three read voltage references: V_{READ2} for lower pages and V_{READ1} and V_{READ3} for upper pages. One could thus evaluate two respective zero/one balances: between the D1 and D2 states for lower pages, and between the E and D1 states combined with between the D2 and D3 states for upper pages. Alternatively, one could carry out three separate lower-page-like reads setting the single lower-page read threshold to values near each of V_{READ1} , V_{READ2} , and V_{READ3} .

By way of example, in some 4LC NVM embodiments, read voltage references V_{READ1} , V_{READ2} , and V_{READ3} may need shifting, at least for one of the two bits stored in the cell. Working similarly to the SLC case, one of the two bits stored in the cell requires a single read at a first read threshold (V_{READ2} , in this MLC case). Ascertaining the other bit requires two additional read thresholds (V_{READ1} , and V_{READ3} , and effectively two associated respective additional reads that are performed internally by the NVM).

According to various embodiments, either: the two read thresholds (V_{READ1} and V_{READ3}) for the other bit are moved in unison by a same and/or a differing amount (under a presumption that drift is similar for both states); or the two read thresholds for the other bit are moved independently (at a cost in additional read operations). The latter option requires knowing the state of the SLC-like bit in the cell (the LSB), as the state of the SLC-like bit determines which of the two read thresholds (V_{READ1} and V_{READ3}) are used for ascertaining the corresponding other bit (the MSB).

In some 4LC embodiments, a combined “upper page” read data zero/one balance is evaluated for a simultaneous blend of the E and D1 distributions and the D2 and D3 distributions. Based on disparities in the combined zero/one balance read vs. that written, and in view of the bit sense reversal (1 to 0 for moving from E to D1, and 0 to 1 for moving from D2 to D3), the corresponding two read thresholds (V_{READ1} and V_{READ3}) are moved in unison in opposite directions. This is done such that the zero/one balance of each component of the blend moves “in the same direction” (e.g., the movements do not conflict).

In some 4LC embodiments, the movement of the two upper page read thresholds (V_{READ1} and V_{READ3}) are separately determined by respective evaluations of zero/one balance for the upper page while using a respective LSB for each bit of the read unit. When the LSB=1, the direction of movement is the opposite of that of the LSB=0 case. For example, instead of computing a disparity between a zero count and a one count in upper page read data, an upper page disparity is computed by selectively modifying each bit of upper page read data based on a corresponding bit of lower page read data, converting a sense of the bits of the upper page read data from a numeric value to a directional value where, for example, 0 represents a higher threshold voltage and 1 represents a lower threshold voltage. In some embodiments, an XNOR of the bits of the upper page read data and the corresponding bits of the lower page read data performs the converting.

Removing Bias from Soft-Decision Samples

In some embodiments, soft decision information is obtained from NVM to perform soft-decision based decoding. The soft decision information is obtained by reading a read unit (or each group of read units transferred at one time from one NVM die) at a variety of read threshold sampling points (for an SLC, values of V_{READ1}) about the nominal (unadjusted) read threshold to obtain multiple samples of data of the read unit, thus building up a soft-decision likelihood of a value of each bit of the read unit. A spacing of the samples depends, at least in part, on particular properties of the NVM used, such as its nominal separation of change states.

The soft-decision information is biased, however, unless the nominal read threshold about which the samples are taken is a read threshold read-equilibrium point (as described above). If the current read threshold is not the threshold read-equilibrium point, then all the soft-decision information is biased in one direction or the other. There are two ways to adjust for this and obtain unbiased (or at least less-biased) soft-decision information:

1. Since a zero/one balance of the read unit as sampled at each of the read threshold sampling points is readily identified, the read threshold read-equilibrium point is readily determined (according to embodiment, such as by using linear interpolation). An offset is calculated corresponding to the difference between the determined read threshold read-equilibrium point and the (old, unadjusted) nominal read threshold. Previously deter-

mined soft-decision information is enabled to be properly “re-centered” and still used by applying the offset with respect to the voltages at which the soft-decision information was previously sampled (with values larger than 1.0 or less than 0.0 being truncated). This approach requires no additional reads, but produces a differing amount of precision/accuracy for some of the data.

2. Since a zero/one balance of the read unit as sampled at each of the read threshold sampling points is readily identified, after collecting all required samples, the read threshold sampling point closest to the read threshold read-equilibrium point is readily determined. The closest is the read threshold sampling point with the zero/one balance most proximate to the zero/one balance written. For the SLC encryption (or scrambling) case, it is the read threshold sampling point closest to having a 50-50 balance. As all of the samples nominally have equivalent spacing from each other, the closest sample is chosen as a new nominal read threshold, and additional soft-decision information samples are optionally gathered (assuming the new nominal read threshold differs from the old). Alternatively, a binary search to find the threshold read-equilibrium point is done first, with precision of the binary search limited to a finest granularity of sampling required for the desired precision of the soft-decision information. The desired precision of the soft-decision information has an associated sample window around the new nominal read threshold. Doing a binary search as a part of multiple reads needed to obtain soft-decision information requires no additional reads, unless the old nominal read threshold falls outside the desired precision sample window.

In the SLC encryption embodiments described so far, the above techniques have focused on finding a preferred read threshold that yields a read data zero/one balance that is closest to 50-50. For cases where the threshold voltage distributions are of generally the same symmetric shape and do not substantively overlap, this preferred read threshold voltage would also correspond to the minima of the combined threshold voltage distributions.

Read Threshold Interpolated from Known Points on the Threshold Voltage Distributions

In some embodiments, another approach to finding the preferred read threshold is to instead find two points, one on each adjacent threshold voltage distribution, and determine the midpoint between these two points by interpolation. For example, sampling at the peak of each threshold voltage distribution should yield a 75/25 balance of read data zeroes/ones (or ones/zeroes, depending on which peak). Once the two peaks are identified, the calculated midpoint between the two on the threshold voltage axis is used to set a new read threshold.

If it is known a priori that the threshold voltage distributions are not uniform (e.g., asymmetrically biased on one side or the other, such as by having a long tail), then in some embodiments that information is used to interpret the position of the peak and to locate the center by a slightly more complex interpolation (e.g., not just the midpoint). Factors such as retention time and/or wear may influence the symmetry of the threshold voltage distributions and are accounted for in some embodiments. Multiple sampling points are also used in some embodiments to reveal the shape of the threshold voltage distributions.

In some embodiments where the range of threshold voltage adjustment is limited, it may not be possible to find the

true peak. In particular, the threshold voltage cannot be moved less than 0V with some flash devices even though the voltage distribution of the E state extends into negative threshold voltages. Knowledge of the threshold voltage distribution still permits the midpoint to be determined via interpolation. For example, if at 0V the read data zero/one balance is 60/40, then approximately 10% excess zeroes are being observed and an area of the E distribution to the right of 0V is approximately 10%. In a first approach of some embodiments, the peak of the D1 distribution is found, and the midpoint is interpolated based on knowledge of the approximate shape of the E distribution. In a second approach of some embodiments, the point of the D1 distribution at which the read data zero/one balance is 40/60 (the opposite point from the 0V measurement) is found. The calculated midpoint between the 60/40 and the 40/60 observations is then used to set the new preferred read threshold. The calculated midpoint could be determined with greater precision with additional knowledge and/or measurement of the two threshold voltage distributions.

In some embodiments the points being interpolated are both on a same side as the read equilibrium point. For example, knowing that a first read threshold sampling point X yields 75/25 read data zero/one balance and a second read threshold sampling point Y yields 62.5/37.5, the read equilibrium point would be close to $Y \pm (X - Y)$ where “+” vs. “-” depends on whether Y is less than X. The direction the read equilibrium point lies is the direction corresponding to moving the read data zero/one balance closer to 50/50. In the example given, that is in the direction from X and toward Y. Knowledge of the threshold voltage distributions would yield a more accurate interpolation rather than the simple linear interpolation indicated.

Example Embodiments

In concluding the introduction to the detailed description, what follows is a collection of example embodiments, including at least some explicitly enumerated as “ECs” (Example Combinations), providing additional description of a variety of embodiment types in accordance with the concepts described herein; these examples are not meant to be mutually exclusive, exhaustive, or restrictive; and the invention is not limited to these example embodiments but rather encompasses all possible modifications and variations within the scope of the issued claims and their equivalents.

EC1) A method, comprising:

determining a disparity between a number of zeroes and a number of ones for at least one or more reads of an NVM; and

in response to an uncorrectable read from the NVM, reducing a subsequent disparity via selectively shifting a threshold voltage based at least in part on the determined disparity.

EC2) The method of EC1, further comprising:

wherein the NVM is an MLC memory having a plurality of states, each of the states having an associated MSB and LSB value in accordance with a predetermined Gray code mapping, the one or more reads of the NVM are MSB reads; and

wherein the determining a disparity comprises using LSB reads corresponding to the one or more MSB reads to selectively invert data of the MSB reads.

EC3) The method of EC1, wherein each of the at least one or more reads is a separate read of at least one portion of a respective page of the NVM.

EC4) The method of EC1, wherein the determining of the disparity comprises determining a difference between the number of zeroes and the number of ones.

13

EC5) The method of EC1, wherein the number of zeroes and the number of ones are determined via respective counts.

EC6) The method of EC4, wherein the shifting is not performed if the absolute value of the difference is below a predetermined amount.

EC7) The method of EC1, wherein the shifting is by a predetermined voltage if the absolute value of the difference is below a predetermined amount.

EC8) The method of EC1, wherein a direction of the threshold voltage shifting is based at least in part on a mapping of logical state to NVM charge state.

EC9) The method of EC1, wherein the NVM is an SLC memory having a mapping of logical one to the erase state, and the threshold voltage is shifted lower based at least in part on a determination that the number of zeroes is fewer than the number of ones.

EC10) The method of EC9, wherein the threshold voltage is shifted lower further based at least in part on a determination that the disparity is in excess of a predetermined amount.

EC11) The method of EC1, wherein the NVM is an SLC memory having a mapping of logical one to the erase state, and the threshold voltage is shifted lower based on a determination that the disparity indicates a predominance of ones.

EC12) The method of EC1, wherein the NVM is an SLC memory having a mapping of logical one to the erase state, and the threshold voltage is shifted higher based at least in part on a determination that the number of zeroes is greater than the number of ones.

EC13) The method of EC12, wherein the disparity is in excess of a predetermined amount.

EC14) The method of EC1, wherein the NVM is an SLC memory having a mapping of logical one to the erase state, and the threshold voltage is shifted higher based on a determination that the disparity indicates a predominance of zeroes.

EC15) The method of EC4, wherein the number of zeroes and the number of ones used in the difference are based on a plurality of reads.

EC16) The method of EC1, wherein the determining is performed at each of a plurality of threshold voltages, and the shifting is based at least in part on the determined disparities.

EC17) A device, comprising:

means for determining a disparity between a number of zeroes and a number of ones for at least one or more reads of an NVM; and

means for, in response to an uncorrectable read from the NVM, reducing a subsequent disparity via selectively shifting a threshold voltage based at least in part on the determined disparity.

EC18) The device of EC17, further comprising:

wherein the NVM is an MLC memory having a plurality of states, each of the states having an associated MSB and LSB value in accordance with a predetermined Gray code mapping, the one or more reads of the NVM are MSB reads; and

wherein the means for determining a disparity comprises means for using LSB reads corresponding to the one or more MSB reads to selectively invert data of the MSB reads.

EC19) The device of EC17, wherein each of the at least one or more reads is a separate read of at least one portion of a respective page of the NVM.

14

EC20) The device of EC17, wherein the determining of the disparity comprises determining a difference between the number of zeroes and the number of ones.

EC21) The device of EC17, wherein the number of zeroes and the number of ones are determined via respective counts.

EC22) The device of EC20, wherein the shifting is not performed if the absolute value of the difference is below a predetermined amount.

EC23) The device of EC20, wherein the number of zeroes and the number of ones used in the difference are based on a plurality of reads.

EC24) The device of EC17, wherein the determining is performed at each of a plurality of threshold voltages, and the shifting is based at least in part on the determined disparities.

EC25) A method, comprising:

determining disparities, each of the disparities having a respective magnitude and being a respective disparity of a respective known distribution of values stored in an NVM and respective values read from the NVM; and

periodically in response to an uncorrectable read from the NVM, selectively shifting at least one read threshold voltage based at least in part on at least one of the disparities determined prior to the shifting.

EC26) The method of EC25, wherein each periodically shifted value of the at least one read threshold voltage is determined in accordance with a goal of keeping below a predetermined amount the respective magnitudes of at least one of the disparities determined subsequent to the shifting.

EC27) The method of EC26, wherein, at least until the next uncorrectable read, the respective magnitudes of the at least one of the disparities determined subsequent to the shifting are reduced below the predetermined amount via a single update of the at least one read threshold voltage.

EC28) The method of EC25, wherein the NVM is an SLC.

EC29) The method of EC25, wherein the NVM is an MLC.

EC30) The method of EC29, wherein the MLC is a 4LC MLC, and the disparities are evaluated with respect to at least two states of an E state and a plurality of D states.

EC31) The method of EC29, wherein the MLC is a 4LC MLC, and the known distribution is 25-25-25-25 percent among four states.

EC32) The method of EC25, wherein the NVM is an MLC having a plurality of states, each of the states having an associated MSB and LSB value in accordance with a predetermined Gray code mapping, the at least one threshold voltage is shifted in view of the determining disparities being performed with respect to lower page reads, and the voltage shifting is in a first direction in response to a predominance of zeroes disparity and in a second direction in response to a predominance of ones disparity.

EC33) The method of EC25, wherein the NVM is an MLC having a plurality of states, each of the states having an associated MSB and LSB value in accordance with a predetermined Gray code mapping, and the at least one threshold voltage is shifted in view of the determining disparities being performed with respect to upper page reads.

EC34) The method of EC33, wherein the upper page reads performed for the determining disparities are performed using a particular LSB for each bit of the read unit, and for a first LSB value the voltage shifting is in a first direction in response to a predominance of zeroes disparity and in a second direction in response to a predominance of

ones disparity, and for a second LSB value the direction sense of the voltage shifting is reversed.

EC35) The method of EC33, wherein the upper page reads performed for the disparity determination are performed using a particular LSB for each bit of the read unit, and the direction of the voltage shifting is determined at least in part based on the particular LSB value of the upper page reads and at least in part on the nature of the disparity.

EC36) The method of EC25, wherein the NVM is an MLC having an E state and a plurality of D states, and the least one threshold voltage is between the E state and a selected one of the plurality of D states.

EC37) The method of EC25, wherein the NVM is an MLC having an E state and a plurality of D states, and the least one threshold voltage is between a selected two adjacent D states of the plurality of D states.

EC38) The method of EC25, further comprising:

wherein the NVM is an MLC memory having a plurality of states, each of the states having an associated MSB and LSB value in accordance with a predetermined Gray code mapping, and for the values read from the NVM, the reads are MSB reads; and

wherein the determining disparities comprises using LSB reads corresponding to the one or more MSB reads to selectively invert data of the MSB reads.

EC39) A device, comprising:

means for determining disparities, each of the disparities having a respective magnitude and being a respective disparity of a respective known distribution of values stored in an NVM and respective values read from the NVM; and

means for, periodically in response to an uncorrectable read from the NVM, selectively shifting at least one read threshold voltage based at least in part on at least one of the disparities determined prior to the shifting.

EC40) The device of EC39, wherein each periodically shifted value of the at least one read threshold voltage is determined in accordance with a goal of keeping below a predetermined amount the respective magnitudes of at least one of the disparities determined subsequent to the shifting.

EC41) The device of EC40, wherein, at least until the next uncorrectable read, the respective magnitudes of the at least one of the disparities determined subsequent to the shifting are reduced below the predetermined, amount via a single update of the at least one read threshold voltage.

EC42) The device of EC39, wherein the NVM is an SLC.

EC43) The device of EC39, wherein the NVM is an MLC.

EC44) A method, comprising:

determining a difference in a zeroes count and a ones count of each of one or more portions of data read from an NVM;

wherein the determining is performed and the difference is calculated at each of a plurality of read reference voltage settings of a read threshold;

using interpolation of the calculated differences to calculate a threshold center-point; and

shifting the read threshold to an updated read reference voltage setting based at least in part on the calculated threshold center-point.

EC45) The method of EC44, wherein the interpolation is a linear interpolation.

EC46) The method of EC44, wherein the determining a difference comprises determining the zeroes count and the ones count for at least one or more reads of the NVM prior to calculating the difference in the zeroes count and the ones count.

EC47) A device, comprising:

means for determining a difference in a zeroes count and a ones count of each of one or more portions of data read from an NVM, wherein the determining is performed and the difference is calculated at each of a plurality of read reference voltage settings of a read threshold;

means for using interpolation of the calculated differences to calculate a threshold center-point; and

means for shifting the read threshold to an updated read reference voltage setting based at least in part on the calculated threshold center-point.

EC48) The device of EC47, wherein the interpolation is a linear interpolation.

EC49) The device of EC47, wherein the means for determining a difference comprises a means for determining the zeroes count and the ones count for at least one or more reads of the NVM prior to calculating the difference in the zeroes count and the ones count.

EC50) A method, comprising:

reading data of a same read unit at each of a plurality of threshold voltages;

for each read of the same read unit, determining a zeroes read count and a ones read count;

calculating a threshold center-point based at least in part on the read counts; and

generating soft-decision information based at least in part on the data read and the calculated threshold center-point.

EC51) The method of EC50, wherein the threshold center-point calculation is performed at least in part using interpolation of calculated differences in the zeroes read count and the ones read count.

EC52) The method of EC51, wherein the interpolation is a linear interpolation.

EC53) The method of EC50, wherein a plurality of values of the soft-decision information are generated based on data read at respective threshold voltage offsets relative to the calculated threshold center-point.

EC54) The method of EC53, wherein an even number of the values are generated corresponding to an equal number of the respective voltage offsets above and below the calculated threshold center-point.

EC55) The method of EC53, wherein use of the calculated threshold center-point as a reference for generating the soft-decision information improves the efficacy of the soft-decision information generated by reducing bias that would otherwise be present.

EC56) A device, comprising:

means for reading data of a same read unit at each of a plurality of threshold voltages;

means for determining a zeroes read count and a ones read count for each read of the same read unit;

means for calculating a threshold center-point based at least in part on the read counts; and

means for generating soft-decision information based at least in part on the data read and the calculated threshold center-point.

EC57) The device of EC56, wherein the threshold center-point calculation is performed at least in part using interpolation of calculated differences in the zeroes read count and the ones read count.

EC58) The device of EC57, wherein the interpolation is a linear interpolation.

EC59) The device of EC56, wherein a plurality of values of the soft-decision information are generated based on data

read at respective threshold voltage offsets relative to the calculated threshold center-point.

EC60) The device of EC59, wherein an even number of the values are generated corresponding to an equal number of the respective voltage offsets above and below the calculated threshold center-point.

EC61) The device of EC59, wherein use of the calculated threshold center-point as a reference for generating the soft-decision information improves the efficacy of the soft-decision information generated by reducing bias that would otherwise be present.

EC62) A method, comprising:

reading data of a same read unit at each of a plurality of threshold voltages;

for each read of the same read unit, determining a zeroes count and a ones count;

based at least in part on the read counts, identifying a closest to center-point threshold voltage of the plurality of threshold voltages, the closest to center-point threshold voltage corresponding to a conceptual threshold center-point; and

generating soft-decision information based at least in part on the data read and the identified closest to center-point threshold voltage.

EC63) The method of EC62, wherein the threshold center-point calculation is performed at least in part using interpolation of calculated differences in the zeroes read count and the ones read count.

EC64) The method of EC63, wherein the interpolation is a linear interpolation.

EC65) The method of EC62, wherein a plurality of values of the soft-decision information are generated based on data read at respective threshold voltage offsets relative to the calculated threshold center-point.

EC66) The method of EC65, wherein a number of the values are generated corresponding to a number of intervals determined by the respective voltage offsets.

EC67) The method of EC62, wherein use of the calculated threshold center-point as a reference for generating the soft-decision information improves the efficacy of the soft-decision information generated by reducing bias that would otherwise be present.

EC68) A device, comprising:

means for reading data of a same read unit at each of a plurality of threshold voltages;

means for determining a zeroes count and a ones count for each read of the same read unit;

means for identifying a closest to center-point threshold voltage of the plurality of threshold voltages, based at least in part on the read counts, the closest to center-point threshold voltage corresponding to a conceptual threshold center-point; and

means for generating soft-decision information based at least in part on the data read and the identified closest to center-point threshold voltage.

EC69) The device of EC68, wherein the threshold center-point calculation is performed at least in part using interpolation of calculated differences in the zeroes read count and the ones read count.

EC70) The device of EC69, wherein the interpolation is a linear interpolation.

EC71) The device of EC68, wherein a plurality of values of the soft-decision information are generated based on data read at respective threshold voltage offsets relative to the calculated threshold center-point.

EC72) The device of EC71, wherein a number of the values are generated corresponding to a number of intervals determined by the respective voltage offsets.

EC73) The device of EC68, wherein use of the calculated threshold center-point as a reference for generating the soft-decision information improves the efficacy of the soft-decision information generated by reducing bias that would otherwise be present.

EC74) A method, comprising:

storing data in an NVM, wherein the stored data has a particular zero-one balance;

maintaining a zeroes count and a ones count, wherein the zeroes count and the ones count are read counts;

evaluating a zero-one balance of the NVM as a function of the zeroes count, the ones count, and the particular zero-one balance;

subsequent to an uncorrectable read from the NVM and a determination that the zero-one balance is outside of a specified range, shifting the threshold voltage to a value determined at least in part based on the evaluating and in accordance with a goal of restoring the zero-one balance to within the specified range in as few shifts of the threshold voltage as possible; and

wherein data recovery from the NVM is optimized.

EC75) The method of EC74, further comprising:

wherein the shifting of the threshold voltage is a shift to a lower voltage based on a determination that the zero-one balance comprises an excess of ones relative to the particular zero-one balance; and

wherein the shifting of the threshold voltage is a shift to a higher voltage based on a determination that the zero-one balance comprises an excess of zeroes relative to the particular zero-one balance.

EC76) The method of EC74, wherein the data is stored as a result of using encryption methods and the particular zero-one balance is 50% zeroes and 50% ones.

EC77) The method of EC74, wherein the particular zero-one balance corresponds to less than 55% zeroes and more than 45% ones.

EC78) The method of EC74, wherein a probability that the particular zero-one balance corresponds to more than 52.5% zeroes is less than 1 in 10,000.

EC79) The method of EC74, wherein the particular zero-one balance corresponds to less than 45% zeroes and more than 55% ones.

EC80) A device, comprising:

means for storing data in an NVM, wherein the stored data has a particular zero-one balance;

means for maintaining a zeroes count and a ones count, wherein the zeroes count and the ones count are read counts;

means for evaluating a zero-one balance of the NVM as a function of the zeroes count, the ones count, and the particular zero-one balance;

means for, subsequent to an uncorrectable read from the NVM and a determination that the zero-one balance is outside of a specified range, shifting the threshold voltage to a value determined at least in part based on the evaluating and in accordance with a goal of restoring the zero-one balance to within the specified range in as few shifts of the threshold voltage as possible; and wherein data recovery from the NVM is optimized.

EC81) The device of EC80, further comprising:

wherein the shifting of the threshold voltage is a shift to a lower voltage based on a determination that the zero-one balance comprises an excess of ones relative to the particular zero-one balance; and

wherein the shifting of the threshold voltage is a shift to a higher voltage based on a determination that the zero-one balance comprises an excess of zeroes relative to the particular zero-one balance.

EC82) The device of EC80, wherein the data is stored as a result of using encryption methods and the particular zero-one balance is 50% zeroes and 50% ones.

EC83) The device of EC80, wherein the particular zero-one balance corresponds to less than 55% zeroes and more than 45% ones.

EC84) The device of EC80, wherein a probability that the particular zero-one balance corresponds to more than 52.5% zeroes is less than 1 in 10,000.

EC85) The device of EC80, wherein the particular zero-one balance corresponds to less than 45% zeroes and more than 55% ones.

EC86) A method, comprising:

scrambling data to be written to a portion of a non-volatile memory to produce scrambled data having a known statistical average number of zero bits and a known statistical average number of one bits;

writing the scrambled data to the portion of the non-volatile memory;

subsequent to the writing, reading the portion of the non-volatile memory;

computing a disparity between a number of zero bits and a number of one bits in data read from the portion of the non-volatile memory; and

determining a threshold voltage for a subsequent read of the non-volatile memory based at least in part on the disparity.

EC87) The method of EC86, wherein the determining is further based on the known statistical average number of zero bits or the known statistical average number of one bits.

EC88) The method of EC86, further comprising:

re-reading the portion of the non-volatile memory using the threshold voltage.

EC89) The method of EC88, wherein the writing comprises encoding the scrambled data with an ECC encoder and writing the encoded scrambled data to the portion of the non-volatile memory, and further comprising:

prior to the re-reading, decoding the data read from the portion of the non-volatile memory with an ECC decoder to determine that the data read from the portion of the non-volatile memory is uncorrectable.

EC90) The method of EC89, wherein the disparity is a first disparity, the threshold voltage is a first threshold voltage, and further comprising:

decoding the data re-read from the portion of the non-volatile memory using the first threshold voltage with the ECC decoder to determine that the data re-read from the portion of the non-volatile memory using the first threshold voltage is uncorrectable;

computing a second disparity between a number of zero bits and a number of one bits in the data re-read from the portion of the non-volatile memory using the first threshold voltage; and

determining a second threshold voltage for a second re-read of the non-volatile memory based at least in part on the first disparity and the second disparity.

EC91) The method of EC90, further comprising:

re-reading the portion of the non-volatile memory using the second threshold voltage.

EC92) The method of EC89, wherein the ECC encoder adds one or more bytes to the scrambled data, and wherein the one or more bytes added by the ECC encoder have the

known statistical average number of zero bits and the known statistical average number of one bits.

EC93) A device, comprising:

means for scrambling data to be written to a portion of a non-volatile memory to produce scrambled data having a known statistical average number of zero bits and a known statistical average number of one bits;

means for writing the scrambled data to the portion of the non-volatile memory;

means for subsequent to the writing, reading the portion of the non-volatile memory;

means for computing a disparity between a number of zero bits and a number of one bits in data read from the portion of the non-volatile memory; and

means for determining a threshold voltage for a subsequent read of the non-volatile memory based at least in part on the disparity.

EC94) The device of EC93, wherein the determining is further based on the known statistical average number of zero bits or the known statistical average number of one bits.

EC95) The device of EC93, further comprising:

means for re-reading the portion of the non-volatile memory using the threshold voltage.

EC96) The device of EC95, wherein the writing comprises encoding the scrambled data with an ECC encoder and writing the encoded scrambled data to the portion of the non-volatile memory, and further comprising:

means for, prior to the re-reading, decoding the data read from the portion of the non-volatile memory with an ECC decoder to determine that the data read from the portion of the non-volatile memory is uncorrectable.

EC97) The device of EC96, wherein the disparity is a first disparity, the threshold voltage is a first threshold voltage, and further comprising:

means for decoding the data re-read from the portion of the non-volatile memory using the first threshold voltage with the ECC decoder to determine that the data re-read from the portion of the non-volatile memory using the first threshold voltage is uncorrectable;

means for computing a second disparity between a number of zero bits and a number of one bits in the data re-read from the portion of the non-volatile memory using the first threshold voltage; and

means for determining a second threshold voltage for a second re-read of the non-volatile memory based at least in part on the first disparity and the second disparity.

EC98) The device of EC97, further comprising:

means for re-reading the portion of the non-volatile memory using the second threshold voltage.

EC99) The device of EC96, wherein the ECC encoder adds one or more bytes to the scrambled data, and wherein the one or more bytes added by the ECC encoder have the known statistical average number of zero bits and the known statistical average number of one bits.

EC100) A tangible computer readable medium having a set of instructions stored therein that when executed by a processing element cause the processing element to perform and/or control operations comprising:

scrambling data to be written to a portion of a non-volatile memory to produce scrambled data having a known statistical average number of zero bits and a known statistical average number of one bits;

writing the scrambled data to the portion of the non-volatile memory;

subsequent to the writing, reading the portion of the non-volatile memory;

computing a disparity between a number of zero bits and a number of one bits in data read from the portion of the non-volatile memory; and

determining a threshold voltage for a subsequent read of the non-volatile memory based at least in part on the disparity.

EC101) The tangible computer readable medium of EC100, wherein the determining is further based on the known statistical average number of zero bits or the known statistical average number of one bits.

EC102) The tangible computer readable medium of EC100, the operations further comprising:

re-reading the portion of the non-volatile memory using the threshold voltage.

EC103) The tangible computer readable medium of EC102, wherein the writing comprises encoding the scrambled data with an ECC encoder and writing the encoded scrambled data to the portion of the non-volatile memory, and the operations further comprise:

prior to the re-reading, decoding the data read from the portion of the non-volatile memory with an ECC decoder to determine that the data read from the portion of the non-volatile memory is uncorrectable.

EC104) The tangible computer readable medium of EC103, wherein the disparity is a first disparity, the threshold voltage is a first threshold voltage, and the operations further comprise:

decoding the data re-read from the portion of the non-volatile memory using the first threshold voltage with the ECC decoder to determine that the data re-read from the portion of the non-volatile memory using the first threshold voltage is uncorrectable;

computing a second disparity between a number of zero bits and a number of one bits in the data re-read from the portion of the non-volatile memory using the first threshold voltage; and

determining a second threshold voltage for a second re-read of the non-volatile memory based at least in part on the first disparity and the second disparity.

EC105) The tangible computer readable medium of EC104, the operations further comprising:

re-reading the portion of the non-volatile memory using the second threshold voltage.

EC106) The tangible computer readable medium of EC103, wherein the ECC encoder adds one or more bytes to the scrambled data, and wherein the one or more bytes added by the ECC encoder have the known statistical average number of zero bits and the known statistical average number of one bits.

System

FIG. 1A illustrates selected details of an embodiment of SSD 101 including an SSD controller using zero/one balance restoration for managing non-volatile storage, such as implemented via NVM elements (e.g., flash memories). SSD Controller 100 is communicatively coupled via one or more External Interfaces 110 to a host (such as host 102 of FIG. 1B). According to various embodiments, External Interfaces 110 are one or more of: a SATA interface; a SAS interface; a PCIe interface; a Fibre Channel interface; an Ethernet Interface (such as 10 Gigabit Ethernet); a non-standard version of any of the preceding interfaces; a custom interface; or any other type of interface used to interconnect storage and/or communications and/or computing devices. For example, in some embodiments, SSD Controller 100 includes a SATA interface and a PCIe interface.

SSD Controller 100 is further communicatively coupled via one or more Device Interfaces 190 to NVM 199 includ-

ing one or more storage devices, such as one or more of Flash Device 192. According to various embodiments, Device Interfaces 190 are one or more of: an asynchronous interface; a synchronous interface; a single-data-rate (SDR) interface; a double-data-rate (DDR) interface; a DRAM-compatible DDR or DDR2 synchronous interface; an ONFI compatible interface, such as an ONFI 2.2 or ONFI 3.0 compatible interface; a Toggle-mode compatible flash interface; a non-standard version of any of the preceding interfaces; a custom interface; or any other type of interface used to connect to storage devices.

Each Flash Device 192 has, in some embodiments, one or more individual Flash Die 194. According to type of a particular one of Flash Device 192, a plurality of Flash Die 194 in the particular Flash Device 192 is optionally and/or selectively accessible in parallel. Flash Device 192 is merely representative of one type of storage device enabled to communicatively couple to SSD Controller 100. In various embodiments, any type of storage device is usable, such as an SLC NAND flash memory, MLC NAND flash memory, NOR flash memory, flash memory using polysilicon or silicon nitride technology-based charge storage cells, two- or three-dimensional technology-based flash memory, read-only memory, static random access memory, dynamic random access memory, ferromagnetic memory, phase-change memory, racetrack memory, ReRAM, or any other type of memory device or storage medium.

According to various embodiments, Device Interfaces 190 are organized as: one or more busses with one or more of Flash Device 192 per bus; one or more groups of busses with one or more of Flash Device 192 per bus, where busses in a group are generally accessed in parallel; or any other organization of one or more of Flash Device 192 onto Device Interfaces 190.

Continuing in FIG. 1A, SSD Controller 100 has one or more modules, such as Host Interfaces 111, Data Processing 121, Buffer 131, Map 141, Recycler 151, ECC 161, Device Interface Logic 191, and CPU 171. The specific modules and interconnections illustrated in FIG. 1A are merely representative of one embodiment, and many arrangements and interconnections of some or all of the modules, as well as additional modules not illustrated, are conceived. In a first example, in some embodiments, there are two or more Host Interfaces 111 to provide dual-porting. In a second example, in some embodiments, Data Processing 121 and/or ECC 161 are combined with Buffer 131. In a third example, in some embodiments, Host Interfaces 111 is directly coupled to Buffer 131, and Data Processing 121 optionally and/or selectively operates on data stored in Buffer 131. In a fourth example, in some embodiments, Device Interface Logic 191 is directly coupled to Buffer 131, and ECC 161 optionally and/or selectively operates on data stored in Buffer 131.

Host Interfaces 111 sends and receives commands and/or data via External Interfaces 110, and, in some embodiments, tracks progress of individual commands via Tag Tracking 113. For example, the commands include a read command specifying an address (such as an LBA) and an amount of data (such as a number of LBA quanta, e.g., sectors) to read; in response the SSD provides read status and/or read data. For another example, the commands include a write command specifying an address (such as an LBA) and an amount of data (such as a number of LBA quanta, e.g., sectors) to write; in response the SSD provides write status and/or requests write data and optionally subsequently provides write status. For yet another example, the commands include a de-allocation command (e.g. a trim command) specifying one or more addresses (such as one or more LBAs) that no

longer need be allocated; in response the SSD modifies the Map accordingly and optionally provides de-allocation status. In some contexts an ATA compatible TRIM command is an exemplary de-allocation command. For yet another example, the commands include a super capacitor test command or a data hardening success query; in response, the SSD provides appropriate status. In some embodiments, Host Interfaces **111** is compatible with a SATA protocol and, using NCQ commands, is enabled to have up to 32 pending commands, each with a unique tag represented as a number from **0** to **31**. In some embodiments, Tag Tracking **113** is enabled to associate an external tag for a command received via External Interfaces **110** with an internal tag used to track the command during processing by SSD Controller **100**.

According to various embodiments, one or more of: Data Processing **121** optionally and/or selectively processes some or all data sent between Buffer **131** and External Interfaces **110**; and Data Processing **121** optionally and/or selectively processes data stored in Buffer **131**. In some embodiments, Data Processing **121** uses one or more Engines **123** to perform one or more of: formatting; reformatting; transcoding; and any other data processing and/or manipulation task.

Buffer **131** stores data sent to/from External Interfaces **110** from/to Device Interfaces **190**. In some embodiments, Buffer **131** additionally stores system data, such as some or all map tables, used by SSD Controller **100** to manage one or more of Flash Device **192**. In various embodiments, Buffer **131** has one or more of: Memory **137** used for temporary storage of data; DMA **133** used to control movement of data to and/or from Buffer **131**; and ECC-X **135** used to provide higher-level error correction and/or redundancy functions; and other data movement and/or manipulation functions. An example of a higher-level redundancy function is a RAID-like capability, where redundancy is at a flash device (e.g., multiple ones of Flash Device **192**) level and/or a flash die (e.g., Flash Die **194**) level instead of at a disk level.

According to various embodiments, one or more of: ECC **161** optionally and/or selectively processes some or all data sent between Buffer **131** and Device Interfaces **190**; and ECC **161** optionally and/or selectively processes data stored in Buffer **131**. In some embodiments, ECC **161** is used to provide lower-level error correction and/or redundancy functions, such as in accordance with one or more ECC techniques. In some embodiments, ECC **161** implements one or more of: a CRC code; a Hamming code; an RS code; a BCH code; an LDPC code; a Viterbi code; a trellis code; a hard-decision code; a soft-decision code; an erasure-based code; any error detecting and/or correcting code; and any combination of the preceding. In some embodiments, ECC **161** includes one or more decoders (such as LDPC decoders).

Device Interface Logic **191** controls instances of Flash Device **192** via Device Interfaces **190**. Device Interface Logic **191** is enabled to send data to/from the instances of Flash Device **192** according to a protocol of Flash Device **192**. Device Interface Logic **191** includes Scheduling **193** to selectively sequence control of the instances of Flash Device **192** via Device Interfaces **190**. For example, in some embodiments, Scheduling **193** is enabled to queue operations to the instances of Flash Device **192**, and to selectively send the operations to individual ones of the instances of Flash Device **192** (or Flash Die **194**) as individual ones of the instances of Flash Device **192** (or Flash Die **194**) are available.

Map **141** converts between data addressing used on External Interfaces **110** and data addressing used on Device

Interfaces **190**, using Table **143** to map external data addresses to locations in NVM **199**. For example, in some embodiments, Map **141** converts LBAs used on External Interfaces **110** to block and/or page addresses targeting one or more Flash Die **194**, via mapping provided by Table **143**. For LBAs that have never been written since drive manufacture or de-allocation, the Map points to a default value to return if the LBAs are read. For example, when processing a de-allocation command, the Map is modified so that entries corresponding to the de-allocated LBAs point to one of the default values. In various embodiments, there are various default values, each having a corresponding pointer. The plurality of default values enables reading some de-allocated LBAs (such as in a first range) as one default value, while reading other de-allocated LBAs (such as in a second range) as another default value. The default values, in various embodiments, are defined by flash memory, hardware, firmware, command and/or primitive arguments and/or parameters, programmable registers, or various combinations thereof.

In some embodiments, Map **141** uses Table **143** to perform and/or to look up translations between addresses used on External Interfaces **110** and data addressing used on Device Interfaces **190**. According to various embodiments, Table **143** is one or more of: a one-level map; a two-level map; a multi-level map; a map cache; a compressed map; any type of mapping from one address space to another, and any combination of the foregoing. According to various embodiments, Table **143** includes one or more of: static random access memory; dynamic random access memory; NVM (such as flash memory); cache memory; on-chip memory; off-chip memory; and any combination of the foregoing.

In some embodiments, Recycler **151** performs garbage collection. For example, in some embodiments, instances of Flash Device **192** contain blocks that must be erased before the blocks are re-writable. Recycler **151** is enabled to determine which portions of the instances of Flash Device **192** are actively in use (e.g., allocated instead of de-allocated), such as by scanning a map maintained by Map **141**, and to make unused (e.g., de-allocated) portions of the instances of Flash Device **192** available for writing by erasing them. In further embodiments, Recycler **151** is enabled to move data stored within instances of Flash Device **192** to make larger contiguous portions of the instances of Flash Device **192** available for writing.

In some embodiments, instances of Flash Device **192** are selectively and/or dynamically configured, managed, and/or used to have one or more bands for storing data of different types and/or properties. A number, arrangement, size, and type of the bands are dynamically changeable. For example, data from a computing host is written into a hot (active) band, while data from Recycler **151** is written into a cold (less active) band. In some usage scenarios, if the computing host writes a long, sequential stream, then a size of the hot band grows, whereas if the computing host does random writes or few writes, then a size of the cold band grows.

CPU **171** controls various portions of SSD Controller **100**. CPU **171** includes CPU Core **172**. CPU Core **172** is, according to various embodiments, one or more single-core or multi-core processors. The individual processors cores in CPU Core **172** are, in some embodiments, multi-threaded. CPU Core **172** includes instruction and/or data caches and/or memories. For example, the instruction memory contains instructions to enable CPU Core **172** to execute programs (e.g. software sometimes called firmware) to control SSD Controller **100**. In some embodiments, some or all

of the firmware executed by CPU Core 172 is stored on instances of Flash Device 192 (as illustrated, e.g., as Firmware 106 of NVM 199 in FIG. 1B).

In Various embodiments, CPU 171 further includes: Command Management 173 to track and control commands received via External Interfaces 110 while the commands are in progress; Buffer Management 175 to control allocation and use of Buffer 131; Translation Management 177 to control Map 141; Coherency Management 179 to control consistency of data addressing and to avoid conflicts such as between external data accesses and recycle data accesses; Device Management 181 to control Device Interface Logic 191; Identity Management 182 to control modification and communication of identify information, and optionally other management units. None, any, or all of the management functions performed by CPU 171 are, according to various embodiments, controlled and/or managed by hardware, by software (such as firmware executing on CPU Core 172 or on a host connected via External Interfaces 110), or any combination thereof.

In some embodiments, CPU 171 is enabled to perform other management tasks, such as one or more of: gathering and/or reporting performance statistics; implementing SMART; controlling power sequencing, controlling and/or monitoring and/or adjusting power consumption; responding to power failures; controlling and/or monitoring and/or adjusting clock rates; and other management tasks.

Various embodiments include a computing-host flash memory controller that is similar to SSD Controller 100 and is compatible with operation with various computing hosts, such as via adaptation of Host Interfaces 111 and/or External Interfaces 110. The various computing hosts include one or any combination of a computer, a workstation computer, a server computer, a storage server, a SAN, a NAS device, a DAS device, a storage appliance, a PC, a laptop computer, a notebook computer, a netbook computer, a tablet device or computer, an ultrabook computer, an electronic reading device (such as an e-reader), a PDA, a navigation system, a (handheld) GPS device, an automotive control system, an automotive media control system or computer, a printer, copier or fax machine or all-in-one device, a POS device, a cash-register, a media player, a television, a media recorder, a DVR, a digital camera, a cellular handset, a cordless telephone handset, and an electronic game.

In various embodiments, all or any portions of an SSD controller (or a computing-host flash memory controller) are implemented on a single IC, a single die of a multi-die IC, a plurality of dice of a multi-die IC, or a plurality of ICs. For example, Buffer 131 is implemented on a same die as other elements of SSD Controller 100. For another example, Buffer 131 is implemented on a different die than other elements of SSD Controller 100.

FIG. 1B illustrates selected details of various embodiments of systems including one or more instances of the SSD of FIG. 1A. SSD 101 includes SSD Controller 100 coupled to NVM 199 via Device Interfaces 190. The figure illustrates various classes of embodiments: a single SSD coupled directly to a host, a plurality of SSDs each respectively coupled directly to a host via respective external interfaces, and one or more SSDs coupled indirectly to a host via various interconnection elements.

As an example embodiment of a single SSD coupled directly to a host, one instance of SSD 101 is coupled directly to Host 102 via External Interfaces 110 (e.g. Switch/Fabric/Intermediate Controller 103 is omitted, bypassed, or passed-through). As an example embodiment of a plurality of SSDs each coupled directly to a host via respective

external interfaces, each of a plurality of instances of SSD 101 is respectively coupled directly to Host 102 via a respective instance of External Interfaces 110 (e.g. Switch/Fabric/Intermediate Controller 103 is omitted, bypassed, or passed-through). As an example embodiment of one or more SSDs coupled indirectly to a host via various interconnection elements, each of one or more instances of SSD 101 is respectively coupled indirectly to Host 102. Each indirect coupling is via a respective instance of External Interfaces 110 coupled to Switch/Fabric/Intermediate Controller 103, and Intermediate Interfaces 104 coupling to Host 102.

Some of the embodiments including Switch/Fabric/Intermediate Controller 103 also include Card Memory 112C coupled via Memory Interface 180 and accessible by the SSDs. In various embodiments, one or more of the SSDs, the Switch/Fabric/Intermediate Controller, and/or the Card Memory are included on a physically identifiable module, card, or pluggable element (e.g. I/O Card 116). In some embodiments, SSD 101 (or variations thereof) corresponds to a SAS drive or a SATA drive that is coupled to an initiator operating as Host 102.

Host 102 is enabled to execute various elements of Host Software 115, such as various combinations of OS 105, Driver 107, Application 109, and Multi-Device Management Software 114. Dotted-arrow 107D is representative of Host Software ↔ I/O Device Communication, e.g. data sent/received to/from one or more of the instances of SSD 101 and from/to any one or more of OS 105 via Driver 107, Driver 107, and Application 109, either via Driver 107, or directly as a VF.

OS 105 includes and/or is enabled to operate with drivers (illustrated conceptually by Driver 107) for interfacing with the SSD. Various versions of Windows (e.g. 95, 98, ME, NT, XP, 2000, Server, Vista, and 7), various versions of Linux (e.g. Red Hat, Debian, and Ubuntu), and various versions of MacOS (e.g. 8, 9 and X) are examples of OS 105. In various embodiments, the drivers are standard and/or generic drivers (sometimes termed “shrink-wrapped” or “pre-installed”) operable with a standard interface and/or protocol such as SATA, AHCI, or NVM Express, or are optionally customized and/or vendor specific to enable use of commands specific to SSD 101. Some drives and/or drivers have pass-through modes to enable application-level programs, such as Application 109 via Optimized NAND Access (sometimes termed ONA) or Direct NAND Access (sometimes termed DNA) techniques, to communicate commands directly to SSD 101, enabling a customized application to use commands specific to SSD 101 even with a generic driver. ONA techniques include one or more of: use of non-standard modifiers (hints); use of vendor-specific commands; communication of non-standard statistics, such as actual NVM usage according to compressibility; and other techniques. DNA techniques include one or more of: use of non-standard commands or vendor-specific providing more direct access to the NVM, such as by bypassing formatting of data that the I/O device would otherwise do; and other techniques. Examples of the driver are a driver without ONA or DNA support, an ONA-enabled driver, a DNA-enabled driver, and an ONA/DNA-enabled driver. Further examples of the driver are a vendor-provided, vendor-developed, and/or vendor-enhanced driver, and a client-provided, client-developed, and/or client-enhanced driver.

Examples of the application-level programs are an application without ONA or DNA support, an ONA-enabled application, a DNA-enabled application, and an ONA/DNA-

enabled application. Dotted-arrow **109D** is representative of Application \leftrightarrow I/O Device Communication (e.g. bypass via a driver or bypass via a VF for an application), e.g. an ONA-enabled application and an ONA-enabled driver communicating with an SSD, such as without the application using the OS as an intermediary. Dotted-arrow **109V** is representative of Application \leftrightarrow I/O Device Communication (e.g. bypass via a VF for an application), e.g. a DNA-enabled application and a DNA-enabled driver communicating with an SSD, such as without the application using the OS or the driver as intermediaries.

One or more portions of NVM **199** are used, in some embodiments, for firmware storage; e.g. Firmware **106**. The firmware storage includes one or more firmware images (or portions thereof). A firmware image has, for example, one or more images of firmware, executed, e.g., by CPU Core **172** of SSD Controller **100**. A firmware image has, for another example, one or more images of constants, parameter values, and NVM device information, referenced, e.g. by the CPU core during the firmware execution. The images of firmware correspond, e.g., to a current firmware image and zero or more previous (with respect to firmware updates) firmware images. In various embodiments, the firmware provides for generic, standard, ONA, and/or DNA operating modes. In some embodiments, one or more of the firmware operating modes are enabled (e.g. one or more APIs are “unlocked”) via keys or various software techniques, optionally communicated and/or provided by a driver.

In some embodiments lacking the Switch/Fabric/Intermediate Controller, the SSD is coupled to the Host directly via External Interfaces **110**. In various embodiments, SSD Controller **100** is coupled to the Host via one or more intermediate levels of other controllers, such as a RAID controller. In some embodiments, SSD **101** (or variations thereof) corresponds to a SAS drive or a SATA drive and Switch/Fabric/Intermediate Controller **103** corresponds to an expander that is in turn coupled to an initiator, or alternatively Switch/Fabric/Intermediate Controller **103** corresponds to a bridge that is indirectly coupled to an initiator via an expander. In some embodiments, Switch/Fabric/Intermediate Controller **103** includes one or more PCIe switches and/or fabrics.

In various embodiments, such as some of the embodiments where Host **102** is a computing host (e.g. a computer, a workstation computer, a server computer, a storage server, a SAN, a NAS device, a DAS device, a storage appliance, a PC, a laptop computer, a notebook computer, and/or a netbook computer), the computing host is optionally enabled to communicate (e.g. via optional I/O & Storage Devices/Resources **117** and optional LAN/WAN **119**) with one or more local and/or remote servers (e.g. optional Servers **118**). The communication enables, for example, local and/or remote access, management, and/or usage of any one or more of SSD **101** elements. In some embodiments, the communication is wholly or partially via Ethernet. In some embodiments, the communication is wholly or partially via Fibre Channel. LAN/WAN **119** is representative, in various embodiments, of one or more Local and/or Wide Area Networks, such as any one or more of a network in a server farm, a network coupling server farms, a metro-area network, and the Internet.

In various embodiments, an SSD controller and/or a computing-host flash memory controller in combination with one or more NVMs are implemented as a non-volatile storage component, such as a USB storage component, a CF storage component, an MMC storage component, an eMMC storage component, a Thunderbolt storage component, a

UFS storage component, an SD storage component, a Memory Stick storage component, and an xD-picture card storage component.

In various embodiments, all or any portions of an SSD controller (or a computing-host flash memory controller), or functions thereof, are implemented in a host that the controller is to be coupled with (e.g., Host **102** of FIG. **1B**). In various embodiments, all or any portions of an SSD controller (or a computing-host flash memory controller), or functions thereof, are implemented via hardware (e.g., logic circuitry), software and/or firmware (e.g., driver software or SSD control firmware), or any combination thereof. For example, functionality of or associated with an ECC unit (such as similar to ECC **161** and/or ECC-X **135** of FIG. **1A**) is implemented partially via software on a host and partially via a combination of firmware and hardware in an SSD controller. For another example, functionality of or associated with a recycler unit (such as similar to Recycler **151** of FIG. **1A**) is implemented partially via software on a host and partially via hardware in a computing-host flash memory controller.

Zero/One Balance Restoration

FIGS. **2A** through **2D** illustrate threshold voltage distributions relevant to “zero/one” balance restoration. For the purposes of this discussion, each hill-like bump is an abstraction for an independent Gaussian-like curve representing a threshold voltage probability distribution for a respective state of a read unit sized portion of an NVM. The threshold voltage axis is drawn with increasing positive voltage to the right. An absolute scale is purposefully not provided, and no reference points are identified, such that the plots apply more generally to a larger population of NVMs. FIGS. **2A** and **2B** are SLC specific, while FIGS. **2C** and **2D** are MLC specific.

In FIG. **2A**, representing an initial time (e.g., when the corresponding data is written), leftmost distribution **220** represents a logical one and the rightmost distribution **230** represents a logical zero. Depending on the technology, the leftmost distribution may reside (at least predominantly) at negative voltages. Read threshold **210** is ideally situated between the two distributions.

Except as linked by the embodiments herein, independent of these threshold voltage distributions, in some embodiments, the NVM is written with specific knowledge of the statistical distribution of zeroes and ones being stored. More particularly, in some embodiments various encryption and/or scrambling techniques are used such that the statistical distribution of zeroes and ones is 50-50 percent (50 percent zeroes and 50 percent ones). When the SLC is read using the nominal read threshold, which is ideally situated for the case of FIG. **2A**, the observed read data statistical distribution of zeroes and ones is likewise 50-50 percent. The statistical distribution being 50-50 percent does not mean that any one sample of data would have an exactly even balance of zero bits and one bits, but rather that an average over many samples produces a ratio of zero bits and one bits that converges on 50-50 percent with an increasingly tight probabilistic bound as a number of the samples increases. An analogy is a distribution of heads and tails when a coin is flipped many times, generating a Gaussian distribution with a variance that is approximated by $n/4$ where n is a number of the coin flips. For example, if there are 18,432 bits in a read unit with a 50-50 percent statistical distribution of zero bits and one bits, the variance in the number of zero bits (or one bits) is approximately 4,608 and the standard deviation is approximately 68. With a standard deviation of 68, less than 1 in one million samples of read units would be

expected to have a number of zero bits more than 340 (5 standard deviations) away from the average of 9,216.

In FIG. 2B, representing a later time, nominal read threshold **210** is as in FIG. 2A, and the two threshold voltage distributions **225** and **235** are shifted with respect to their earlier respective distributions **220** and **230** in FIG. 2A. For the purposes of example, the two distributions are shown as having both uniformly shifted to the left (toward more negative voltages). It should be understood that more generally the two distributions are enabled to move independently of each other and in either a positive or negative direction.

In view of FIG. 2B, when the SLC is again read using the nominal read threshold **210**, it is predictable that the observed statistical distribution of zeroes and ones read directly from the NVM (e.g., prior to any error correction) will not be 50-50 percent. More particularly, for the conceptual example given, an erroneous excess of ones is to be expected, as the read threshold is such that some of the zeroes will be falsely read as ones.

In practice, the direction of inference is reversed. That is, in practice, generally such shifts in the threshold voltage distributions are not known or directly knowable. In some embodiments, instead the observation of a disparity in the zeroes and ones read from the NVM (with respect to the known stored distribution of zeroes and ones) is used to infer the existence of shifts in the threshold voltage distributions. Furthermore, in these embodiments the read threshold is adjusted based at least in part on the disparity observed (as detailed elsewhere herein) to read threshold **215**, as required until the zero/one balance is restored.

Similarly, FIG. 2C represents an initial time (such as when the corresponding data is written) with initial threshold voltage distributions, while FIG. 2D represents a later time with corresponding later threshold voltage distributions. More specifically, for a selected Gray code mapping, respectively representing the 11, 10, 00, and 10 states, threshold voltage distributions **251**, **261**, **271**, and **281** (of FIG. 2D) are shifted with respect to their earlier respective threshold voltage distributions **250**, **260**, **270**, and **280** (of FIG. 2C). Three nominal (initial) read thresholds are also shown: V_{READ1} **240**, V_{READ2} **242**, and V_{READ3} **244**. Again for the purposes of example, in FIG. 2D the four distributions are shown as having all uniformly shifted to the left (toward more negative voltages). It should be understood that more generally the four distributions are enabled to move independently of each other and in either a positive or negative direction.

In some embodiments, various scrambling techniques are used such that the statistical distribution of the four states is 25-25-25-25 percent (25 percent in each state). When the MLC is read using the nominal read thresholds, which are ideally situated for the case of FIG. 2C, in some embodiments the statistical distribution of the four states is enabled to be confirmed to likewise be 25-25-25-25 percent. (As discussed elsewhere, zeroes and ones are directly observable by the controller, but all four states are inferable as required.) In some embodiments, the observation of a disparity (a deviation from the expected 25-25-25-25 percent) read from the NVM (with respect to the known stored distribution of states) is used to infer the existence of shifts in the threshold voltage distributions. The read thresholds are then adjusted (as detailed elsewhere herein) as shown in FIG. 2D to become V_{READ1} **265**, V_{READ2} **275**, and V_{READ3} **285**. In some embodiments, the read threshold adjustments are performed

separately for lower pages reads (adjusting V_{READ2} **275**) and upper pages reads (adjusting V_{READ1} **265** and/or V_{READ3} **285**).

Functions for Zero/One Balance Restoration

FIG. 3A provides alternate views of SSD **101**, SSD Controller **100**, and NVM **199** of FIGS. 1A and 1B, redrawing selected logic block boundaries to emphasize particular functions. SSD **101** at the top level features External Interface **110**, SSD Controller **100**, NVM **199**, and Device Interface **190**.

External Interface **110** is coupled to a host, such as Host **102** of FIG. 1B, and supports a high-level storage protocol such as SATA, comprising host-sourced storage-related commands and write-data and controller-sourced read-data, and as detailed in the above discussion of FIG. 1A. Device Interface **190** supports low-level NVM I/O transactions, as detailed in the discussion of FIG. 1A. NVM **199** features NVM Array **332**, Programmable Read Voltage Circuitry **334**, Control/Status Registers **336**, and I/O **338**. SSD Controller **100** features Host Interface **111**, Scrambler **310**, ECC Encoder **320**, Device Interface Logic **191**, Imbalance Detector **340**, Buffer **131**, ECC Decoder **360**, Descrambler **370**, and Balance Restoration Logic **380**. Host Interface **111** couples with the host via External Interface **110** as discussed previously, provides Write Data **306** to Scrambler **310**, and receives Read Data **351** from Buffer **131**.

The write path and functionality is as follows. Scrambler **310** operates on Write Data **306** to generate Scrambled Write Data **311** to ECC Encoder **320**. Scrambler **310** scrambles Write Data **306** in a reversible manner and such that the Scrambled Write Data **311** has a known statistical distribution of states stored. In some embodiments, Scrambler **310** performs encryption, such as via an AES encoder, which randomizes the data as a by-product of encryption. In some embodiments Scrambler **310** uses a Linear Feedback Shift Register (LFSR) to randomize the data (but without any data security intent). ECC Encoder **320** processes the Scrambled Write Data **311** to add additional ECC bits, resulting in ECC Encoded Write Data **321**, which is provided to Device Interface **191** for storing into NVM **199** via Device Interface **190**.

The basic read path and functionality is as follows. NVM pages, comprising Uncorrected and Scrambled (Raw) Read Data **359**, are received from NVM **199** via Device Interface **190**. In some embodiments and as detailed in FIG. 3B, each NVM page read comprises four codewords (a.k.a. read units): Codeword 0 through Codeword 3. Continuing with respect to FIG. 3A, each codeword nominally comprises the data scrambled by Scrambler **310** as well as the additional ECC bits generated by ECC Encoder **320**, but as generally unintentionally altered as a result of NVM storage and retrieval operations. The codewords are provided to ECC Decoder **360** via Buffer **131**. ECC Decoder **360** processes the codewords, exploiting the included additional ECC bits to generally correct any errors that may have arisen at any point and time since the data was first ECC encoded, resulting in Corrected Read Data **361**, which is provided to Descrambler **370**. The Descrambler **370** operates on the Corrected Read Data **361**, reversing the scrambling performed by Scrambler **310**, resulting in Unscrambled Read Data **371**, which is provided as Read Data **351** to Host Interface **111** via Buffer **131**.

The read path and functionality further includes Imbalance Detector **340** and Balance Restoration Logic **380**. Imbalance Detector **340** monitors the Raw Read Data **359** received from the NVM, and at least temporarily maintains data on the statistical distributions of state. Balance Resto-

ration Logic **380** interacts with NVM Control/Status Registers **336**, via Device Interface Logic **191** and NVM I/O **338**, to selectively adjust NVM read thresholds (used to read the NVM Array **332**) via NVM Programmable Read Voltage Circuitry **334**. The read thresholds are adjusted based upon
 5 a) an Uncorrectable Read **362** detected by ECC Decoder **360**, and b) an Imbalance **341** (in the statistical of state) detected by Imbalance Detector **340**. The read unit is binary data in both SLC and MLC embodiments. Disparities in the zeroes and ones are directly measurable in the SLC case (or
 10 when an MLC is operated as though it were an SLC). In some embodiments, disparities in the more than two states of an MLC are inferred based on knowledge of how the data was stored in the NVM.

FIG. **3A** illustrates functions for zero/one balance restoration in a context where a host operates as a source and a sink for data written to and read from NVM **199**. In various embodiments (such as illustrated by FIG. **1A** and FIG. **1B**), one or more other agents operate as sources and sinks for data written to and read from NVM **199**. An example of the
 20 other agents is Recycler **151** of FIG. **1A**, that moves data from one location to another in NVM **199**, as described elsewhere herein.

In various embodiments, any or all portions of functionality associated with any one or more of Scrambler **310**, ECC Encoder **320**, ECC Decoder **360**, and Descrambler **370** are implemented at least in part using one or more of Data Processing **121** and ECC **161** (both of FIG. **1A**).

Control Flows for Zero/One Balance Restoration

FIG. **4** illustrates selected control flow details for an embodiment of the SSD of FIGS. **1A**, **1B**, and **3A**, providing particular details regarding write related operations and read related operations. Write operations start with action **410** and continue through **430**. While a return path is not explicitly drawn, subsequent write operations start anew with action
 30 **410**. Similarly, read operations start with action **440**, continue through **460**, and conditionally continue through **480**. While a return path is not explicitly drawn, subsequent read operations start anew with action **440**. It is assumed that at least an initial write precedes a read to the same location. Otherwise, except for contrived cases, such as in testing, the read operations and write operations are generally invoked independently as application needs dictate, for the same or generally different locations. Nevertheless, in a heuristic example, the first read operation **440** conceptually continues
 35 immediately after the last write operation **430** for a same location.

Considering in more detail the write related operations, original data to be written to the NVM is generally first scrambled in action **410**, such as via encryption or use of an LFSR. Except where more specificity is employed, the term “scrambling” as used herein refers to the manipulation of the original data to be written via any manipulation or selection that results in a known statistical distribution among the various states stored in the NVM. Scrambling methods include encryption and LFSR techniques. (See also Scrambler **310** of FIG. **3A**.) In some encryption embodiments, the scrambling is in accordance with at least one version of the AES encryptions specification. In some LFSR embodiments, the scrambling is in accordance with at least one choice of LFSR. In some embodiments, the scrambling is more directly determined via the selection of one or more test patterns, providing a chosen statistical distribution of states when written to the NVM (without necessarily using encryption or LFSR processing).

The LFSR is a hardware shift register having a combi- national logic feedback network generating an input bit from

a linear function of selected taps of the shift register, or a software model thereof. The LFSR generates a deterministic pseudorandom sequence. At least conceptually, the LFSR-generated pseudorandom sequence is modulo 2 added to the original data to provide the scrambled data.
 5

Subsequently, the scrambled data is next generally (but optionally) ECC encoded, as illustrated in action **420**. (See also ECC Encoder **320** of FIG. **3A**.) ECC encoding provides redundant information that facilitates data recovery (via error correction) in the face of errors that arise for a variety of reasons when using NVM. Such errors are more prevalent, and thus ECC encoding more advisable, with one or more of smaller geometries, greater use, greater age, greater temperature excursions, and use of MLC memories.

The scrambled and optionally ECC encoded data is then written to the NVM, as shown in action **430**, the last of the write related operations. The distribution of states being written is known as a result of the scrambling. In the case of the use of SLC memories and AES encryption, the zero-one distribution is known to be 50-50 (i.e., 50 percent zeroes and 50 percent ones). In the case of the use of a 4LC (4-level MLC) memories and AES encryption, the distribution after writing both lower and upper pages is known to be 25-25-25-25 (i.e., 25 percent in each of the four states).
 20

In some embodiments, at least one reference region is dedicated for the use of one or more test patterns as discussed above. In some embodiments the location of the reference region is selected to minimize (or conversely, to maximize) the impact on the test patterns of various NVM write/read artifacts. In some embodiments, the reference region is a reference page, and its location is chosen to be a last page written in a block, so as to minimize the impact of write disturbs on the data written. As some of the phenomena affecting the voltage distributions are time and/or temperature variant, the reference region serves as a known measure of these effects on zero/one disparity.
 25

Considering in more detail the read related operations, starting in action **440**, one or more read units are read from the NVM and the distribution of states is captured. In some embodiments, each of a zeroes count and a ones count are directly determined. Next, in action **450**, an attempt is made to recover the original data. This performs the inverse of whatever data manipulation operations were performed prior to writing. Generally, ECC encoding was added prior to writing, and thus generally the read units next undergo ECC decoding. (See also ECC Decoder **360** of FIG. **3A**.) On a per read unit basis, if ECC decoding is successful, then descrambling is performed next, as appropriate. E.g., where AES encryption was used prior to writing, AES-decryption is performed. Or, where an LFSR-generated pseudorandom sequence was used, the same sequence is modulo 2 added to the ECC decoded data from the read unit.
 30

The read related operations conclude early when the original data is recoverable, either without error or for errors that are correctable via ECC decoding. See the “no” path from decision **460**. However, on uncorrectable errors of any of the read units, actions **470** and **480** are also performed. See the “yes” path from decision **460**. In action **470**, a disparity in the distribution of states is evaluated. (See also Imbalance Detector **340** of FIG. **3A**.) In some embodiments the determined disparity is a determined difference between the zeroes count and the ones count. In various embodiments, the determined disparity is computed over all of the read units, even if less than all of the read units had uncorrectable errors.
 35

Next, in action **480** an adjusted value (a.k.a. a threshold shift) for at least one read threshold is determined at least in

part based on the magnitude of the determined disparity. In some embodiments, if the magnitude of the determined disparity is below a predetermined tolerance, the read threshold is not adjusted. In some embodiments a multi-factor “correction algorithm” is employed for determining the adjusted read threshold, for which the magnitude of the determined disparity is but one factor.

In some embodiments, the correction algorithm takes the form of a lookup table used to select a replacement value for the read threshold. The magnitude of the determined disparity is the basis for at least part of the index in the lookup table. In some embodiments, the lookup table is pre-characterized (weighted) based on the particular vendor of the NVM being used. In some embodiments, the index has additional components, comprising one or more of: P/E cycles, age, retention time (time since the last write), temperature, and any other factors that may pertain to drift of, or disturbances to, the threshold voltages distributions of the NVM. (See also Balance Restoration Logic **380** of FIG. **3A**.)

Here and throughout the detailed description, it should be remembered that in commonly used flash memory micro-architectures, generally the granularity of data being accessed is different between reads and writes. The minimum size of a read access (a read unit) is a page. The page is organized as one or more codewords, four in some embodiments, as discussed elsewhere herein. In contrast, the minimum size of a write access is a block of many pages.

In some embodiments, disparities are evaluated on an entire page basis. In some embodiments, disparities are evaluated on a codeword basis. The disparity evaluation performed in conjunction with an uncorrectable error on a read, is performed with respect to the same (page, or codeword) location as written in the most recent block write including that same location.

Example Implementation Techniques

In some embodiments, various combinations of all or portions of operations performed by an zero-one balanced SSD, e.g., with flash memories, a computing-host flash memory controller, and/or an SSD controller (such as SSD Controller **100** of FIG. **1A**), and portions of a processor, microprocessor, system-on-a-chip, application-specific-integrated-circuit, hardware accelerator, or other circuitry providing all or portions of the aforementioned operations, are specified by a specification compatible with processing by a computer system. The specification is in accordance with various descriptions, such as hardware description languages, circuit descriptions, netlist descriptions, mask descriptions, or layout descriptions. Example descriptions include: Verilog, VHDL, SPICE, SPICE variants such as PSpice, IBIS, LEF, DEF, GDS-II, OASIS, or other descriptions. In various embodiments, the processing includes any combination of interpretation, compilation, simulation, and synthesis to produce, to verify, or to specify logic and/or circuitry suitable for inclusion on one or more integrated circuits. Each integrated circuit, according to various embodiments, is designable and/or manufacturable according to a variety of techniques. The techniques include a programmable technique (such as a field or mask programmable gate array integrated circuit), a semi-custom technique (such as a wholly or partially cell-based integrated circuit), and a full-custom technique (such as an integrated circuit that is substantially specialized), any combination thereof, or any other technique compatible with design and/or manufacturing of integrated circuits.

In some embodiments, various combinations of all or portions of operations as described by a computer readable medium having a set of instructions stored therein, are

performed by execution and/or interpretation of one or more program instructions, by interpretation and/or compiling of one or more source and/or script language statements, or by execution of binary instructions produced by compiling, translating, and/or interpreting information expressed in programming and/or scripting language statements. The statements are compatible with any standard programming or scripting language (such as C, C++, Fortran, Pascal, Ada, Java, VBscript, and Shell). One or more of the program instructions, the language statements, or the binary instructions, are optionally stored on one or more computer readable storage medium elements. In various embodiments some, all, or various portions of the program instructions are realized as one or more functions, routines, sub-routines, in-line routines, procedures, macros, or portions thereof.

CONCLUSION

Certain choices have been made in the description merely for convenience in preparing the text and drawings and unless there is an indication to the contrary the choices should not be construed per se as conveying additional information regarding structure or operation of the embodiments described. Examples of the choices include: the particular organization or assignment of the designations used for the figure numbering and the particular organization or assignment of the element identifiers (the callouts or numerical designators, e.g.) used to identify and reference the features and elements of the embodiments.

The words “includes” or “including” are specifically intended to be construed as abstractions describing logical sets of open-ended scope and are not meant to convey physical containment unless explicitly followed by the word “within.”

Although the foregoing embodiments have been described in some detail for purposes of clarity of description and understanding, the invention is not limited to the details provided. There are many embodiments of the invention. The disclosed embodiments are exemplary and not restrictive.

It will be understood that many variations in construction, arrangement, and use are possible consistent with the description, and are within the scope of the claims of the issued patent. For example, interconnect and function-unit bit-widths, clock speeds, and the type of technology used are variable according to various embodiments in each component block. The names given to interconnect and logic are merely exemplary, and should not be construed as limiting the concepts described. The order and arrangement of flow-chart and flow diagram process, action, and function elements are variable according to various embodiments. Also, unless specifically stated to the contrary, value ranges specified, maximum and minimum values used, or other particular specifications (such as flash memory technology types; and the number of entries or stages in registers and buffers), are merely those of the described embodiments, are expected to track improvements and changes in implementation technology, and should not be construed as limitations.

Functionally equivalent techniques known in the art are employable instead of those described to implement various components, sub-systems, operations, functions, routines, sub-routines, in-line routines, procedures, macros, or portions thereof. It is also understood that many functional aspects of embodiments are realizable selectively in either hardware (e.g., generally dedicated circuitry) or software (e.g., via some manner of programmed controller or processor), as a function of embodiment dependent design con-

straints and technology trends of faster processing (facilitating migration of functions previously in hardware into software) and higher integration density (facilitating migration of functions previously in software into hardware). Specific variations in various embodiments include, but are not limited to: differences in partitioning; different form factors and configurations; use of different operating systems and other system software; use of different interface standards, network protocols, or communication links; and other variations to be expected when implementing the concepts described herein in accordance with the unique engineering and business constraints of a particular application.

The embodiments have been described with detail and environmental context well beyond that required for a minimal implementation of many aspects of the embodiments described. Those of ordinary skill in the art will recognize that some embodiments omit disclosed components or features without altering the basic cooperation among the remaining elements. It is thus understood that much of the details disclosed are not required to implement various aspects of the embodiments described. To the extent that the remaining elements are distinguishable from the prior art, components and features that are omitted are not limiting on the concepts described herein.

All such variations in design are insubstantial changes over the teachings conveyed by the described embodiments. It is also understood that the embodiments described herein have broad applicability to other computing and networking applications, and are not limited to the particular application or industry of the described embodiments. The invention is thus to be construed as including all possible modifications and variations encompassed within the scope of the claims of the issued patent.

What is claimed is:

1. A storage system comprising:

a non-volatile memory (NVM) comprising a storage device; and

a storage controller communicatively connected with the storage device via a device interface, the storage controller configured to produce scrambled data,

encode the scrambled data by adding at least one byte to the scrambled data to produce encoded scrambled data, the at least one byte having a statistical average number of zero bits and a statistical average number of one bits,

write the encoded scrambled data to a portion of the NVM,

perform a data read of a portion of the NVM using a first threshold voltage,

decode data read from the portion of the NVM to determine that the data read from the portion of the NVM is uncorrectable, and

responsive to an uncorrectable data read from the NVM,

determine a disparity between a nominal statistical distribution of zero bits and one bits in data written to the portion of the NVM and a monitored statistical distribution of zero bits and one bits in the data read from the portion of the NVM, resulting in a determined disparity, and

determine a second threshold voltage, different from the first threshold voltage, for a subsequent read of the portion of the NVM based at least in part on the determined disparity;

whereby re-reading the portion of the NVM at the second threshold voltage requires a smaller number of reads to find an optimal point at which to attempt hard-decision decoding than are required when reading the portion of the NVM using the first threshold voltage.

2. The storage system of claim 1, wherein the storage controller is further configured to encode the scrambled data with an error-correcting code (ECC) encoder, and to decode the data read from the portion of the NVM with an ECC decoder.

3. The storage system of claim 1, wherein the storage controller is further configured to, subsequent to performing the data read of the portion of the NVM, monitor the statistical distribution of the zero bits and the one bits in the data read, and at least temporarily store the monitored statistical distribution of the zero bits and the one bits.

4. The storage system of claim 1, wherein the re-reading of the portion of the NVM at the second threshold voltage results in a subsequent disparity that is less than the determined disparity.

5. The storage system of claim 2, wherein the ECC encoder is configured to provide error correction and redundancy functions comprising an ECC technique.

6. The storage system of claim 5, wherein the ECC technique comprises at least one of a cyclic redundancy check (CRC) code, a Hamming code, a Reed-Solomon (RS) code, a Bose Chaudhuri Hocquenghem (BCH) code, a low-density parity-check (LDPC) code, a Viterbi code, a trellis code, a hard-decision code, a soft-decision code, and an erasure-based code.

7. The storage system of claim 1, wherein the scrambled data has a particular zero-one balance.

8. A method comprising:

storing scrambled data in a non-volatile memory (NVM); encoding the scrambled data by adding at least one byte to the scrambled data, the at least one byte having a statistical average number of zero bits and a statistical average number of one bits;

writing the encoded scrambled data to a portion of the NVM;

performing a data read of the portion of the NVM using a first threshold voltage;

decoding data read from the portion of the NVM to determine that the data read from the portion of the NVM is uncorrectable; and

responsive to an uncorrectable data read from the NVM, determining a disparity between a nominal statistical distribution of zero bits and one bits in data written to the portion of the NVM and a monitored statistical distribution of zero bits and one bits in the data read from the portion of the NVM, resulting in a determined disparity, and

determining a second threshold voltage, different from the first threshold voltage, for a subsequent read of the portion of the NVM based at least in part on the determined disparity;

whereby re-reading the portion of the NVM at the second threshold voltage requires a smaller number of reads to find an optimal point at which to attempt hard-decision decoding than are required when reading the portion of the NVM using the first threshold voltage.

9. The method of claim 8, wherein the encoding is performed with an error-correcting code (ECC) encoder, and the decoding is performed with an ECC decoder.

10. The method of claim 8, further comprising, subsequent to performing the data read of the portion of the NVM:

37

monitoring the statistical distribution of the zero bits and the one bits in the data read; and
at least temporarily storing the monitored statistical distribution of the zero bits and the one bits.

11. The method of claim 8, wherein the re-reading of the portion of the NVM at the second threshold voltage results in a subsequent disparity that is less than the determined disparity.

12. The method of claim 9, wherein the ECC encoder is configured to provide error correction and redundancy functions comprising an ECC technique.

13. The method of claim 12, wherein the ECC technique comprises at least one of a cyclic redundancy check (CRC) code, a Hamming code, a Reed-Solomon (RS) code, a Bose Chaudhuri Hocquenghem (BCH) code, an low-density parity-check (LDPC) code, a Viterbi code, a trellis code, a hard-decision code, a soft-decision code, and an erasure-based code.

14. The method of claim 9, wherein the ECC decoder comprises a low-density parity-check (LDPC) decoder.

15. A storage controller, the storage controller configured to:

perform a data read from a portion of a non-volatile memory (NVM) using a first threshold voltage;

decode data read from the portion of the NVM to determine that the data read from the portion of the NVM is uncorrectable; and

responsive to an uncorrectable data read from the NVM, determine a disparity between a nominal statistical distribution of zero bits and one bits in data written to the portion of the NVM and a monitored statistical

38

distribution of zero bits and one bits in the data read from the portion of the NVM, resulting in a determined disparity, and

determine a second threshold voltage, different from the first threshold voltage, for a subsequent read of the portion of the NVM based at least in part on the determined disparity;

whereby re-reading the portion of the NVM at the second threshold voltage requires a smaller number of reads to find an optimal point at which to attempt hard-decision decoding than are required when reading the portion of the NVM using the first threshold voltage.

16. The storage controller of claim 15, wherein the storage controller is further configured to encode data with an error-correcting code (ECC) encoder, and to decode the data read from the portion of the NVM with an ECC decoder.

17. The storage controller of claim 15, wherein the storage controller is further configured to, subsequent to performing the data read of the portion of the NVM, monitor the statistical distribution of the zero bits and the one bits in the data read, and at least temporarily store the monitored statistical distribution of the zero bits and the one bits.

18. The storage controller of claim 15, wherein the re-reading of the portion of the NVM at the second threshold voltage results in a subsequent disparity that is less than the determined disparity.

19. The storage controller of claim 16, wherein the ECC encoder is configured to provide error correction and redundancy functions comprising an ECC technique.

20. The storage controller of claim 15, wherein the storage controller is further configured to produce scrambled data using a Linear Feedback Shift Register (LFSR).

* * * * *